COLLEGES SCIENCES
BRETAGNE POUR L'INGENIEUR
LOIRE ET LE NUMERIQUE

IMT Atlantique
Bretagne-Pays de la Loire
École Mines-Télécom

THÈSE DE DOCTORAT DE

L'ÉCOLE NATIONALE SUPÉRIEURE
MINES-TÉLÉCOM ATLANTIQUE BRETAGNE PAYS DE LA LOIRE –
IMT ATLANTIQUE

ECOLE DOCTORALE N° 648
Sciences pour l'Ingénieur et le Numérique
Spécialité : *Génie industriel ; productique ; automatique et robotique*

Par

Yuqi LIU

Data-Driven Predictive Control for Building Thermal Management: A Hierarchical Framework Comparing Indirect and Direct Approaches

Thèse présentée et soutenue à IMT Atlantique, Nantes, le 04 / 05 / 2026
Unité de recherche : LS2N UMR 6004
Thèse N° : 2026IMTA0536

Rapporteurs avant soutenance :

Thierry Poinot	Professeur, ENSIP, Poitiers
Guillaume Sandou	Professeur, CentraleSupelec, Gif/Yvette

Composition du Jury :

Président :	Guillaume Sandou	Professeur, CentraleSupelec, Gif/Yvette
Examinateurs :	Stéphane Lecoeuche	Professeur, IMT Mines Alès
	Ionela Prodan	Maitre de conférence HDR, ESISAR, Valence
	Thierry Poinot	Professeur, ENSIP, Poitiers
Dir. de thèse :	Philippe Chevrel	Professeur, IMT-Atlantique, Nantes
Co-encadrants :	Fabien Claveau	Maître assistant, IMT-Atlantique, Nantes
	Pauline Kergus	Chargé de recherche, CNRS, Toulouse

Invité(s)

Hervé Illy	Chargé de recherche	CSTB, Marne-la-Vallée
Gautier Avril	Directeur technique	Société Purecontrol, Rennes

Acknowledgement [1]

This work was funded by the Pays de la Loire Region and the company Purecontrol, under the supervision of Dr. Philippe Chevrel, Dr. Fabien Claveau (Codex team, IMT Atlantique), and Dr. Pauline Kergus (LAPLACE Laboratory), and supported by the SPIN Doctoral School.

The dissertation defense was successfully held on May 4, 2026, leading to the award of the Ph.D. degree. The author would like to express sincere gratitude to all the aforementioned contributors for their valuable support and involvement. The author would also like to express sincere gratitude to the reviewers and members of the defense committee for their careful evaluation of this dissertation and for their valuable remarks and constructive suggestions.

Remerciements

Ce travail a été financé par la Région Pays de la Loire et l'entreprise Purecontrol, sous la direction de Dr Philippe Chevrel, Dr Fabien Claveau (équipe Codex, IMT Atlantique) et Dr Pauline Kergus (laboratoire LAPLACE), avec le soutien de l'École doctorale SPIN.La soutenance de cette thèse, tenue le 4 mai 2026, a conduit à l'obtention du grade de docteur. L'auteur adresse ses sincères remerciements à l'ensemble des personnes et institutions mentionnées ci-dessus pour leur soutien et leur contribution à ce travail. Il remercie également les rapporteurs et les membres du jury pour le temps consacré à l'évaluation de cette thèse.

L'auteur a eu la chance de passer sept années d'études en France, au cours desquelles il a bénéficié de l'accompagnement de nombreux enseignants, chercheurs et encadrants. Ces expériences ont constitué un fondement solide pour son parcours professionnel et personnel.

Il souhaite exprimer une gratitude particulière à ses trois directeurs et encadrants pour l'aide précieuse qu'ils lui ont apportée, tant sur le plan académique que personnel. Au-delà des conseils scientifiques, ils lui ont transmis l'esprit d'initiative, la persévérance et l'exigence professionnelle, des qualités qui représentent pour lui une richesse essentielle de la formation en sciences de l'ingénieur en France.

L'auteur tient également à remercier Dr Alexandre Hache pour les échanges enrichissants et l'expérience partagée tout au long du travail de recherche, ainsi que l'ensemble de ses collègues pour leur ouverture d'esprit et les discussions stimulantes qui ont contribué à créer un environnement de travail inspirant.

致 谢

本研究工作得到了法国卢瓦尔河地区（Pays de la Loire）及 Purecontrol 公司的资助，由 Philippe Chevrel 博士、Fabien Claveau 博士（IMT Atlantique Codex 团队）以及 Pauline Kergus 博士（LAPLACE 实验室）共同指导，并由 SPIN 博士生院负责管理与支持。论文已于二〇二六年五月四日顺利完成答辩，并获授博士学位。在此，谨向上述所有机构与个人致以诚挚的感谢，感谢他们在本研究过程中给予的支持与帮助。作者同时感谢各位论文评审专家和答辩委员会成员对本论文提出的宝贵意见与建议。

作者在法国度过了七年的学习与科研时光，其间得到了众多老师、研究人员和同事的指导与帮助。这段经历不仅为作者的职业发展奠定了坚实基础，也成为人生中宝贵的财富。特别感谢三位导师在科研工作和日常生活中给予的关怀与支持。他们不仅提供了严谨而富有启发性的学术指导，更以自身的专业精神与人格力量，使作者学会了开拓创新、在困难中保持耐心与韧性，并始终以积极的态度面对挑战。

此外，作者还要感谢家人与朋友长期以来给予的关心、陪伴与鼓励，正是他们的支持使作者得以坚持并完成这段求学旅程。“不忘初心，方得始终。”在未来的职业生涯中，作者希望能够以自身所学与积累的经验，继续促进中法两国在科学研究与工程技术领域的交流与合作。

[1] The English, French, and Chinese versions of acknowledgment may contain slight variations resulting from linguistic and stylistic adaptations.

Abstract

A significant share of global energy consumption is attributable to buildings, particularly for heating and cooling services. Model Predictive Control (MPC) is a promising approach for reducing this consumption while ensuring occupants' thermal comfort. However, large-scale deployment faces two major challenges: the strong heterogeneity of building equipment and the limited availability of operational data, which are often constrained by comfort requirements. This dissertation investigates data-driven predictive control approaches to address these challenges by extracting system behavior from historical data.

The work proposes a hierarchical control framework that decouples high-level thermal planning from low-level equipment regulation. Low-level controllers (inner-loop control) are tuned directly from operational data using Virtual Reference Feedback Tuning (VRFT), without requiring explicit equipment models. At the supervisory level, physics-informed RC models identified via least-squares estimation are compared with black-box models obtained using subspace identification methods. Validation on the BOPTEST benchmark shows that RC-MPC achieves substantial energy savings, even with limited data, as the physical structure constrains predictions to thermodynamically feasible trajectories. In contrast, Subspace Predictive Control (SPC), while exhibiting comparable open-loop accuracy, shows inferior closed-loop performance.

Model-free predictive control, referred to as Data-enabled Predictive Control (DeePC), generates predictions directly from input–output trajectories organized in Hankel matrices, without an explicit model identification step. Unlike SPC, which relies on fixed models identified offline, DeePC constructs online dynamic representations through regularized combinations of measured trajectories. Benchmarking results demonstrate that DeePC outperforms both RC-MPC and SPC by simultaneously reducing energy consumption and thermal discomfort, at the cost of more delicate regularization tuning and increased computational requirements.

To account for time-varying dynamics induced by seasonal weather conditions and equipment aging, adaptive extensions of the proposed approaches are developed. Adaptive SPC updates state-space models through recursive identification, whereas adaptive DeePC directly updates the Hankel matrices. Two strategies are investigated: a sliding-window approach incorporating data quality assessment, and a recursive update scheme based on a forgetting factor. Validation results indicate that these adaptive methods yield substantially improved energy–comfort trade-offs under varying operating conditions.

Résumé

Une part significative de l'énergie mondiale est consommée par les bâtiments, en particulier pour les services de chauffage et de climatisation. La MPC classique constitue une approche prometteuse pour réduire cette consommation tout en garantissant le confort thermique des occupants. Toutefois, son déploiement à grande échelle se heurte à deux défis majeurs : d'une part, la forte hétérogénéité des équipements d'un bâtiment à l'autre et, d'autre part, la disponibilité limitée de données opérationnelles. En pratique, les bâtiments sont généralement équipés d'un système de régulation en boucle fermée avant toute collecte de données, et celle-ci se fait typiquement en exploitation normale, sous contraintes strictes de confort et de sécurité. Cette thèse explore des approches de contrôle prédictif fondé sur les données afin de relever ces défis en extrayant le comportement du système à partir de l'historique des données.

Architecture de contrôle hiérarchique

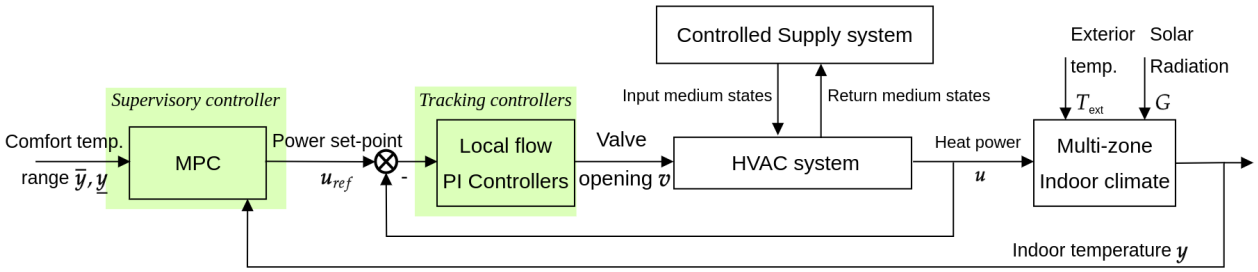

Figure A1: Architecture de commande hiérarchique avec un contrôleur prédictif de supervision et des contrôleurs de suivi de puissance agissant au plus proche des moyens d'action (vannes des radiateurs). Le contrôleur superviseur effectue des prédictions sur l'horizon H avec un pas de temps T_{sup}, en mettant à jour les consignes de puissance thermique à chaque instant. Entre deux mises à jour du superviseur (c'est-à-dire pendant chaque période T_{sup}), les régulateurs PI discrets assurent le suivi des consignes de puissance avec un pas de temps T_{track}.

Le système de contrôle thermique des bâtiments se compose de plusieurs éléments, comme illustré sur la Figure A1. Les signaux de température, de puissance et de vanne sont des vecteurs pour plusieurs zones. Au niveau haut, le contrôleur superviseur reçoit les plages de température de confort souhaitées et génère des références de puissance thermique. Ces références sont ensuite transmises aux contrôleurs de niveau bas, qui ajustent les actionneurs du système CVC (Chauffage, Ventilation et Climatisation), dit aussi HVAC, pour maintenir les conditions souhaitées. Le système CVC délivre la puissance thermique aux zones du bâtiment, dont les températures sont influencées par ailleurs par les conditions extérieures (température ambiante, rayonnement solaire) et les perturbations internes. Cette boucle fermée vise à d'assurer le respect en température des plages de confort tout en optimisant la consommation énergétique.

La structure hiérarchique présente plusieurs avantages pour la gestion thermique des bâtiments. Le contrôleur superviseur opère avec un pas de temps plus long (typiquement 15 à 30 minutes), ce qui permet d'anticiper les évolutions thermiques sur un horizon de plusieurs heures et d'exploiter la flexibilité offerte par l'inertie thermique du bâtiment. Les contrôleurs de niveau bas, quant à eux, fonctionnent avec un pas de temps plus court (quelques minutes) pour assurer un suivi rapide des références de puissance et compenser les dynamiques non-linéaires des équipements. Cette séparation des échelles de temps réduit la complexité du problème d'optimisation au niveau superviseur tout en garantissant une régulation précise des actionneurs.

Méthodologie proposée

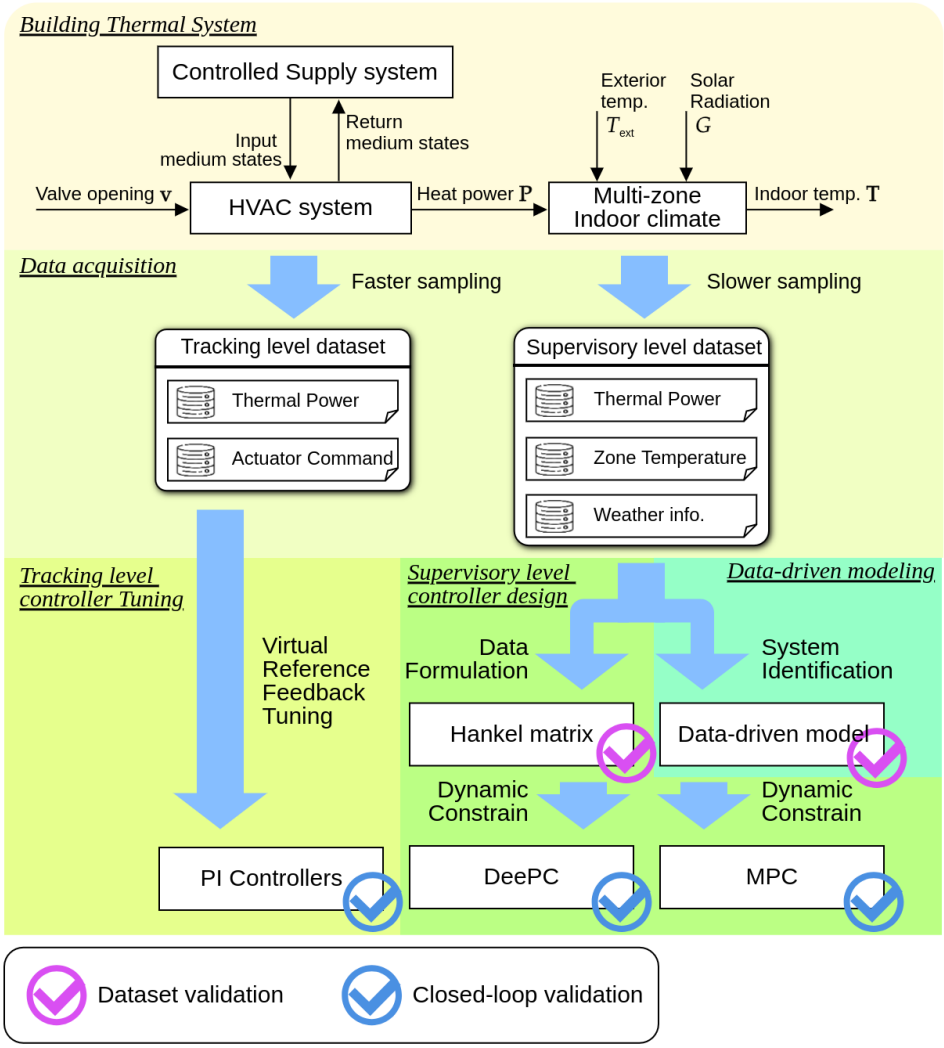

Figure A2: Méthodologie de mise en œuvre proposée. L'acquisition de données collecte des jeux de données à différentes fréquences d'échantillonnage pour les niveaux de suivi et superviseur. La conception des contrôleurs se déroule indépendamment pour les deux niveaux. La validation s'effectue à deux niveaux : évaluation des modèles de prédiction sur jeux de données et tests en boucle fermée des contrôleurs complets.

La mise en œuvre de l'architecture de contrôle hiérarchique suit un processus systématique illustré par la Figure A2, allant de l'acquisition de données à la validation des contrôleurs. Cette méthodologie sépare le développement du niveau de suivi (tracking) et du niveau superviseur, chaque niveau ayant des exigences distinctes en termes de données, de procédures de conception et d'étapes de validation.

L'acquisition de données s'effectue à deux échelles temporelles distinctes. Au niveau de suivi, un échantillonnage rapide capture les dynamiques des équipements CVC (puissance thermique et actionneurs) à travers des expériences d'excitation dédiées utilisant des signaux binaires pseudo-aléatoires. Cette approche expérimentale est pratique car elle peut être conduite sans compromettre le confort des occupants. Au niveau superviseur, un échantillonnage plus lent, aligné sur les dynamiques thermiques du bâtiment, collecte les températures de zones, les puissances délivrées et les perturbations météorologiques, soit à partir de l'exploitation normale, soit à travers d'expériences thermiques spécifiques.

Le réglage des contrôleurs de niveau bas utilise la méthode VRFT (Virtual Reference Feedback Tuning) pour optimiser les paramètres des régulateurs PI directement à partir des données, sans nécessiter de modèle explicite du système. La validation en boucle fermée vérifie ensuite que les contrôleurs peuvent suivre les changements de consigne de puissance dans les délais requis, garantissant ainsi que les références du niveau superviseur peuvent être fidèlement réalisées par les actionneurs CVC.

Au niveau superviseur, deux voies conceptuelles sont explorées, dites indirectes (via un modèle explicite) ou directement à partir des données. D'une part les méthodes basées sur des modèles (RC-MPC utilisant des réseaux RC équivalents, et SPC utilisant l'identification par sous-espaces 4SID), et d'autre part, l'approche sans modèle explicite DeePC (Data-enabled Predictive Control) qui encode le comportement du système dans des matrices de Hankel construites à partir de trajectoires historiques. La validation des jeux de données évalue la qualité des modèles ou formulations à travers des tests de prédiction, tandis que la validation en boucle fermée analyse les performances multi-objectifs via des fronts de Pareto, examine la robustesse aux variations d'hyperparamètres, et évalue l'efficacité computationnelle. Ces différentes approches (VRFT, RC-MPC, SPC, DeePC) et leurs extensions adaptatives seront détaillées dans les sections suivantes.

Dynamique thermique du bâtiment

La gestion thermique des bâtiments multi-zones représente un défi particulier en raison de la diversité des espaces et des systèmes de chauffage. Les zones, caractérisées par les taille et l'orientation des fenêtres, la nature des matériaux de construction et les profils d'occupation, présentent des comportements thermiques distincts. Chaque zone présente ainsi des caractéristiques distinctes en termes de taille, d'équipements CVC, et de sollicitations thermiques. Les Figures A3 et A4 illustrent deux configurations typiques : un bâtiment résidentiel équipé d'un système hydronique avec radiateurs, et un bâtiment de bureaux équipé d'un système de ventilation. Ces différences architecturales et technologiques influencent directement les stratégies de contrôle à adopter, notamment en termes de vitesse de réponse des équipements et de couplage thermique entre zones adjacentes.

Le comportement thermique d'un bâtiment présente des caractéristiques particulières qui influencent directement la conception des stratégies de contrôle. La nature multi-zones du bâtiment implique que chaque espace (bureau, salle de réunion, couloir) possède sa propre dynamique thermique, déterminée par sa capacité thermique, son isolation, et ses échanges avec les zones adjacentes et l'extérieur. Ces zones sont couplées thermiquement, ce qui signifie que le chauffage ou le refroidissement d'une zone affecte les zones voisines par conduction à travers les parois internes.

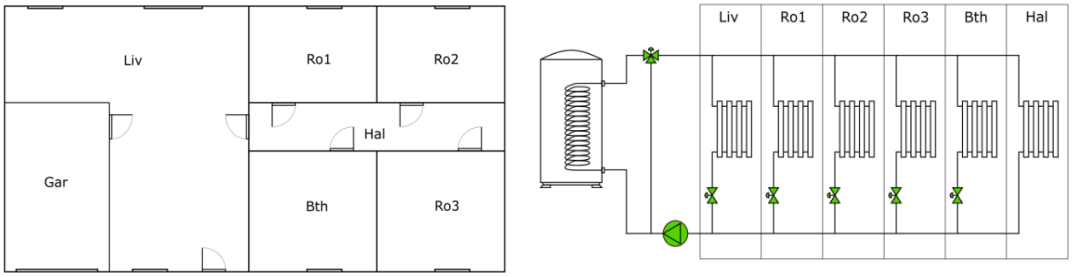

Figure A3: BOPTEST benchmark Multi-zone Residential Hydronic (MZRH) test case.

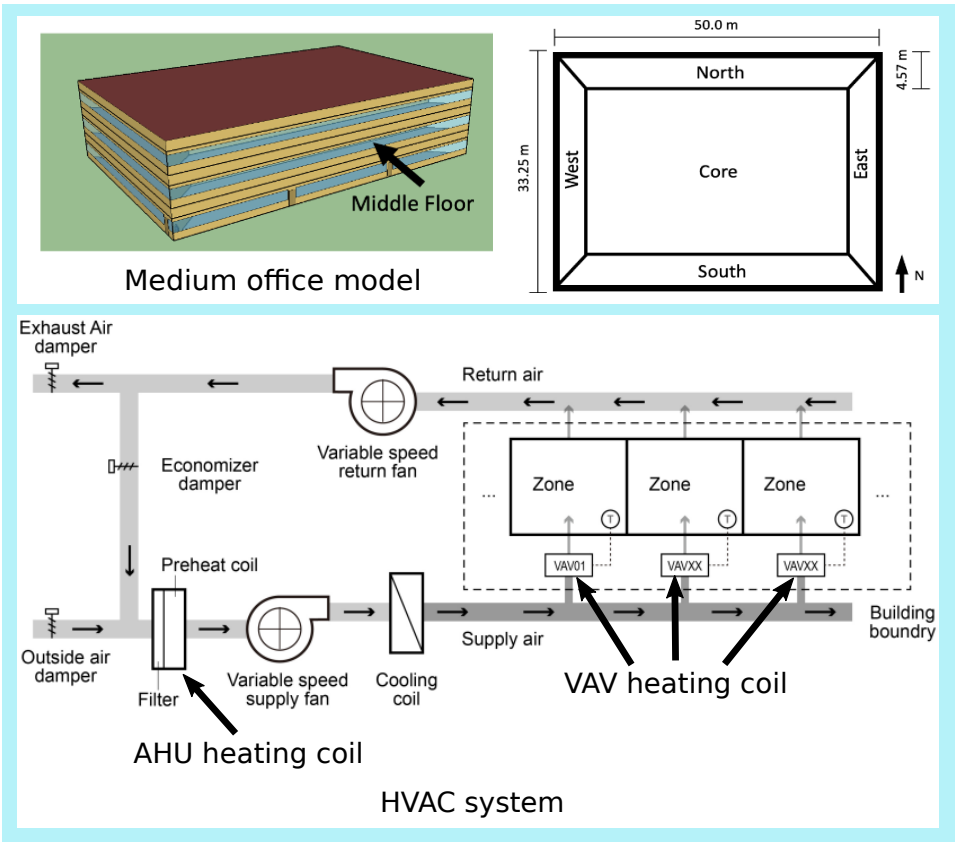

Figure A4: BOPTEST benchmark Multi-zone Office Simple Air (MOSA) test case.

Le système thermique est également soumis à diverses perturbations qui varient sur différentes échelles de temps. Les conditions météorologiques extérieures (température ambiante, rayonnement solaire) suivent des cycles quotidiens et saisonniers prévisibles, tandis que les gains internes dus à l'occupation, à l'éclairage et aux équipements électriques présentent des variations plus stochastiques. Le rayonnement solaire, en particulier, peut contribuer significativement aux apports thermiques gratuits mais varie en fonction de l'orientation des façades, de la nébulosité, et de la saison. Ces perturbations doivent être anticipées par le contrôleur prédictif pour exploiter les apports gratuits et minimiser la consommation.

L'inertie thermique du bâtiment constitue une ressource essentielle pour le contrôle prédictif. Les masses thermiques (murs, dalles, mobilier) accumulent et restituent l'énergie thermique avec

des constantes de temps de plusieurs heures. Cette propriété permet de décaler temporellement les besoins de chauffage ou de refroidissement, par exemple en préchauffant le bâtiment pendant les heures creuses où l'électricité est moins coûteuse, ou en profitant des apports solaires matinaux pour réduire les besoins en fin de journée.

Contrôle prédictif par modèle et modélisation thermique

Le contrôle prédictif par modèle repose sur l'utilisation d'un modèle mathématique du système pour prédire son comportement futur et calculer la séquence de commandes optimales. Dans le contexte de la gestion thermique des bâtiments, la MPC résout à chaque instant un problème d'optimisation qui minimise un critère combinant la consommation d'énergie et les violations de confort sur un horizon de prédiction, tout en respectant les contraintes physiques du système (limites de puissance, plages de température).

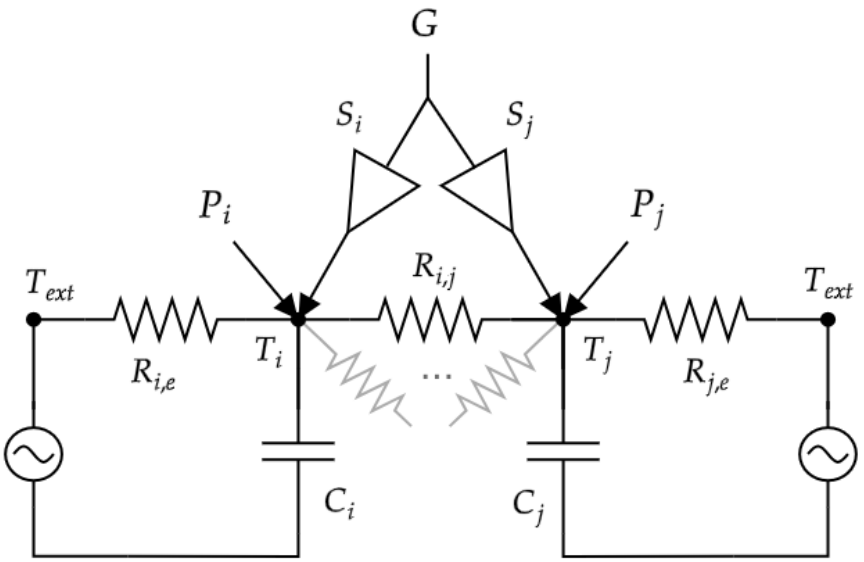

Figure A5: Modèle RC multi-zones du premier ordre. Chaque zone i est caractérisée par une capacité thermique C_i et des résistances thermiques $R_{i,e}$ (vers l'extérieur) et $R_{i,j}$ (vers les zones adjacentes). Les sources de puissance P_i représentent les systèmes de chauffage, tandis que S_i et S_j modélisent les apports solaires.

Deux approches de modélisation sont comparées dans ce travail. La première repose sur des modèles physiques de type RC (Résistance-Capacité), illustrés à la Figure A5, qui représentent le comportement thermique du bâtiment par analogie avec des circuits électriques. Chaque zone est modélisée par une capacité thermique (représentant l'inertie) connectée à l'extérieur et aux zones adjacentes par des résistances thermiques (représentant l'isolation). Les apports solaires et les puissances de chauffage sont modélisés comme des sources de courant. La structure physique de ces modèles garantit que les prédictions restent cohérentes avec les lois de la thermodynamique, ce qui limite les erreurs de prédiction même avec des données d'identification limitées.

Les paramètres des modèles RC (résistances, capacités) sont identifiés à partir de données opérationnelles par la méthode des moindres carrés. Cette approche tire parti de la linéarité du modèle

pour obtenir une solution analytique du problème d'identification, ce qui la rend particulièrement robuste et rapide. De plus, la structure physique du modèle permet d'incorporer des connaissances a priori (par exemple, les ordres de grandeur typiques des résistances thermiques) et de garantir l'interprétabilité des paramètres identifiés.

La seconde approche de modélisation repose sur les méthodes d'identification par sous-espaces (subspace identification), qui génèrent des modèles en représentation d'état sans imposer de structure physique particulière. Ces modèles boîte noire sont capables de capturer des dynamiques complexes mais nécessitent généralement plus de données pour atteindre une précision similaire. Les résultats de validation sur le benchmark BOPTEST montrent que les modèles RC, malgré leur simplicité, permettent d'obtenir des prédictions de qualité suffisante pour le contrôle prédictif, en particulier lorsque la durée des campagnes de mesure est limitée. Le RC-MPC démontre des économies d'énergie substantielles tout en maintenant le confort thermique.

En revanche, le Subspace Predictive Control (SPC), bien qu'affichant une précision comparable en simulation en boucle ouverte, présente des performances inférieures en boucle fermée. Cette dégradation s'explique par le compromis biais-variance inhérent aux méthodes d'identification par sous-espaces : la flexibilité de ces modèles (un atout en boucle ouverte, où ils capturent des dynamiques complexes sans a priori physique) se transforme en sensibilité accrue aux incertitudes en régime rétroactif. Deux mécanismes principaux en sont responsables : l'accumulation des erreurs de prédiction au fil des itérations de contrôle, amplifiée par la rétroaction en boucle fermée, et la sensibilité des modèles boîte noire aux conditions opérationnelles différentes de celles rencontrées lors de l'identification (par exemple, perturbations non modélisées ou dynamiques non linéaires). La contrainte physique imposée par la structure RC limite cette dérive en confinant les prédictions à des trajectoires thermodynamiquement plausibles.

Régulation de niveau bas par VRFT

La performance du contrôle prédictif superviseur dépend fortement de la capacité des contrôleurs de niveau bas à suivre fidèlement les références de puissance thermique. Les équipements CVC présentent des dynamiques non-linéaires et des délais de réponse qui, s'ils ne sont pas correctement compensés, peuvent dégrader significativement les performances globales du système. Par exemple, comme illustré sur la Figure A6, lorsqu'une référence de puissance diminue brutalement, la puissance thermique réellement délivrée décroît progressivement en raison de l'inertie du système de distribution (temps de mélange dans les conduites, constantes de temps des radiateurs). Si le modèle superviseur ne prend pas en compte cette dynamique, il peut générer des références irréalistes et provoquer des violations de confort.

Le problème de contrôle au niveau hiérarchique bas est formulé comme un problème multivariable de suivi de puissance pour les systèmes CVC hydrauliques équipés de plusieurs radiateurs. Le système est modélisé comme un processus à entrées et sorties multiples (MIMO) où le vecteur d'entrée représente les signaux de commande discrets et le vecteur de sortie correspond aux puissances thermiques effectivement délivrées. En raison des interactions hydrauliques et thermiques au sein du réseau CVC, les puissances délivrées sont intrinsèquement couplées, même avec une implémentation décentralisée des contrôleurs.

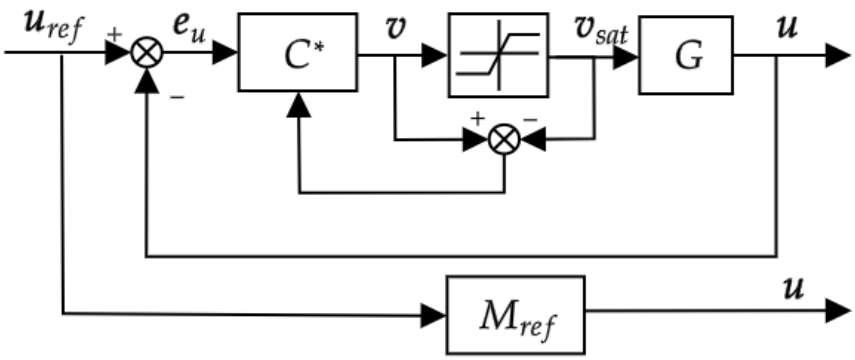

Figure A6: Diagramme bloc du processus VRFT. C^* représente un contrôleur idéal qui pilote la boucle fermée pour atteindre la performance désirée M_{ref}.

Le scénario typique du contrôle de suivi de puissance se présente comme suit :

- Le système CVC est un système à entrées et sorties multiples avec un vecteur de données d'entrée (ouvertures de vannes) et un vecteur de données de sortie (puissances thermiques délivrées).

- Un ensemble de contrôleurs discrets est attendu pour atteindre l'objectif de suivi de puissance. La structure du contrôleur est choisie comme un PID discret avec compensation anti-windup, incluant les gains proportionnel, intégral et dérivé, ainsi qu'un gain anti-windup pour gérer la saturation des actionneurs.

- Les modèles de référence représentent la réponse désirée de la boucle fermée interne, spécifiant le comportement visé pour chaque canal de puissance.

La Figure A6 illustre la méthode VRFT. Cette approche optimise une classe de contrôleurs donnée (PID avec anti-windup) pour approximer un contrôleur idéal en utilisant uniquement les données d'entrée-sortie du système. Le contrôleur idéal est défini comme le contrôleur qui, s'il était inséré dans la boucle fermée durant l'expérience de collecte de données, aurait fourni exactement le comportement désiré. Le VRFT génère une référence virtuelle en simulant à rebours la performance désirée de la boucle fermée. Cette référence virtuelle permet au VRFT de construire les données d'entrée nécessaires au contrôleur. En utilisant cette référence virtuelle et la sortie réelle du système, le VRFT ajuste les paramètres du contrôleur pour correspondre à la performance du contrôleur idéal sans nécessiter de modèle explicite du système.

La modélisation première des couplages hydrauliques et des dynamiques d'actionneurs étant complexe, incertaine et variable dans le temps, cette approche data-driven est particulièrement intéressante pour gérer la grande variété d'équipements rencontrés dans les systèmes de bâtiments. Un modèle de référence diagonal est sélectionné pour spécifier le comportement désiré en boucle fermée de chaque canal de puissance :

- Chaque élément diagonal est défini comme une fonction de transfert du premier ordre stable à gain statique unitaire, à phase minimale, augmentée d'un zéro rapide dont la constante de temps est d'un ordre de grandeur inférieur à celle du pôle dominant.

- Ce choix préserve l'inversibilité et le conditionnement numérique requis par le VRFT, tout en garantissant des temps de réponse maîtrisés sans dépassement et en limitant les erreurs de suivi en régime permanent.

Bien que le modèle de référence soit diagonal, le réglage VRFT est effectué dans un cadre multivariable en utilisant simultanément les données de puissance mesurées sur tous les radiateurs. Les contrôleurs apprennent ainsi implicitement les effets de couplage du système CVC malgré une implémentation décentralisée, représentant un compromis délibéré entre performance et complexité où le comportement local visé reste explicite tandis que les interactions inter-canaux sont gérées par robustesse. L'inclusion de l'action anti-windup est essentielle pour éviter un réglage biaisé en présence de saturation des actionneurs, inévitable dans les systèmes de chauffage à puissance limitée.

L'intégration du VRFT dans l'architecture hiérarchique améliore la cohérence entre les prédictions du superviseur et le comportement réel du système. En calibrant les contrôleurs PI pour reproduire une dynamique de boucle fermée connue, le niveau haut peut incorporer un modèle simplifié de cette dynamique dans ses prédictions, ce qui réduit les erreurs de modélisation et améliore la performance globale. Les résultats expérimentaux confirment que cette approche permet de réduire les écarts entre les références de puissance et celles effectivement délivrées, ce qui se traduit par un meilleur suivi des consignes de température et une réduction de la consommation d'énergie.

Contrôle prédictif sans modèle : Data-enabled Predictive Control(DeePC)

La DeePC représente une approche alternative qui génère directement des prédictions à partir de trajectoires d'entrée-sortie mesurées, sans étape explicite d'identification de modèle. Le principe repose sur le théorème fondamental de Willems, qui établit que tout comportement futur d'un système linéaire peut être représenté comme une combinaison linéaire de trajectoires passées observées, à condition que les trajectoires recueillies soient suffisamment riches en information (persistance d'excitation).

La formulation d'optimisation de DeePC est

$$\min_{g} \quad J + \lambda_y ||\sigma_y||_2^2 + \lambda_g ||g||_p \tag{1a}$$

$$\text{s.t.} \begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} u_{ini} \\ y_{ini} \\ u \\ y \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_y \\ 0 \\ 0 \end{bmatrix}, \tag{1b}$$

$$u \in \mathscr{U}^N, \tag{1c}$$

$$\text{avec} \begin{bmatrix} U_p \\ U_f \end{bmatrix} := \mathscr{H}_{T_{ini}+N}(u_d), \begin{bmatrix} Y_p \\ Y_f \end{bmatrix} := \mathscr{H}_{T_{ini}+N}(y_d) \tag{1d}$$

où J est l'objectif économique. u et y sont les trajectoires de prédiction d'entrée et de sortie. N est le nombre de pas sur l'horizon de prédiction H. $\mathscr{U}^N$ est l'ensemble des contraintes d'entrée sur N pas de prédiction. y_{ini} fait référence à la séquence de retour la plus récente. u_{ini} est la séquence d'entrée la plus récente appliquée au système. T_{ini} est la longueur de y_{ini} et u_{ini}. u_d et y_d désignent les données d'entrée et de sortie d'un ensemble de données d'entraînement informatif qui est persistamment excitant d'ordre $T_{ini} + N$. U_p se compose des T_{ini} premières lignes de $\mathscr{H}_{T_{ini}+N}(u_d)$ pour l'estimation de la condition initiale. U_f se compose des N dernières lignes pour la prédiction de trajectoire future. De même pour Y_p et Y_f. g est un mappeur de trajectoire, indexant les trajectoires historiques dans les

matrices de Hankel pour une combinaison linéaire. σ_y est une variable d'écart qui assure la faisabilité. λ_y et λ_g sont les pondérations sur la contrainte souple et la régularisation de g respectivement.

La prédiction de DeePC repose sur la matrice de Hankel $\mathscr{H}_L(x)$, définie par (2) pour un exemple de série temporelle $x := \{x(t)\}_{t=1}^{T}$ avec un temps d'échantillonnage T_{DeePC}.

$$\mathscr{H}_L(x) := \begin{bmatrix} x(1) & x(2) & \dots & x(T-L+1) \\ x(2) & x(3) & \dots & x(T-L+2) \\ \vdots & \vdots & \ddots & \vdots \\ x(L) & x(L+1) & \dots & x(T) \end{bmatrix}, \tag{2}$$

où L est la hauteur de la matrice.

Concrètement, la DeePC organise les données historiques dans des matrices de Hankel qui empilent les séquences d'entrées et de sorties sur des fenêtres temporelles glissantes. À chaque instant de contrôle, la DeePC résout un problème d'optimisation qui cherche les coefficients de combinaison linéaire des trajectoires historiques permettant de prédire le comportement futur du système, tout en minimisant un critère de coût énergétique et de confort. Des termes de régularisation sont ajoutés pour assurer la stabilité numérique et pour pénaliser l'utilisation de trajectoires peu pertinentes ou corrompues par le bruit.

La DeePC présente plusieurs avantages par rapport aux approches basées sur des modèles. Premièrement, elle évite les erreurs structurelles liées à la nature supposée du modèle choisi. Deuxièmement, elle peut s'adapter naturellement à des changements dans le comportement du système en mettant à jour les matrices de Hankel avec de nouvelles données. Troisièmement, elle peut incorporer des contraintes complexes et non-linéaires directement dans le problème d'optimisation, sans nécessiter de linéarisation ou d'approximation.

Néanmoins, la DeePC présente également des défis importants liés à son paramétrage. Le problème d'optimisation du DeePC consiste à trouver les coefficients de combinaison linéaire des trajectoires historiques qui minimisent le critère de coût énergétique et de confort, tout en satisfaisant les contraintes du système. Ce problème fait intervenir plusieurs hyperparamètres critiques dont le réglage conditionne la stabilité et la performance du contrôle.

Le réglage des poids de régularisation est crucial pour équilibrer la fidélité aux données historiques et la robustesse face au bruit de mesure. Les termes de régularisation, typiquement de type ℓ_1 ou ℓ_2, pénalisent l'utilisation excessive de trajectoires peu pertinentes ou corrompues, garantissant ainsi la stabilité numérique du problème d'optimisation.

Le choix de la profondeur de l'état initial T_{ini} (voir équation (1) dans chapitre 3, le nombre de pas de temps passés utilisés pour estimer l'état actuel du système) présente un compromis délicat entre qualité de prédiction et efficacité computationnelle. D'une part, une grande valeur de T_{ini} permet d'utiliser plus de données historiques pour caractériser l'état actuel, ce qui améliore la précision des prédictions en capturant mieux les dynamiques du système et en réduisant l'influence des perturbations instantanées. D'autre part, une grande valeur de T_{ini} augmente significativement la taille des matrices de Hankel, ce qui accroît le temps de calcul nécessaire pour résoudre le problème d'optimisation à chaque pas de contrôle.

Le problème d'optimisation de la DeePC (équation 1) est convexe et quadratique. Le critère économique J incorpore des objectifs de confort et d'énergie quadratiques par rapport aux variables de décision. Cette convexité garantit la stabilité numérique et l'optimalité globale de la solution, permettant l'utilisation de solveurs quadratiques efficaces pour une résolution fiable en temps réel.

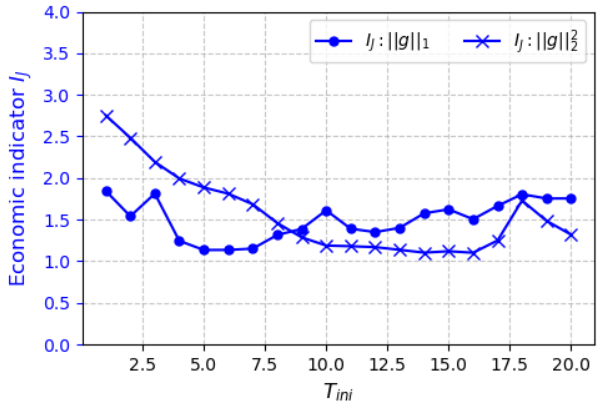

Figure A7: Sensibilité de la performance économique à l'ordre d'estimation de l'état initial T_{ini}. Les disques bleus correspondent à la régularisation ℓ_1 et les croix bleues à la régularisation ℓ_2. Poids de régularisation : $\lambda_y = 1$, $\lambda_g = 100$.

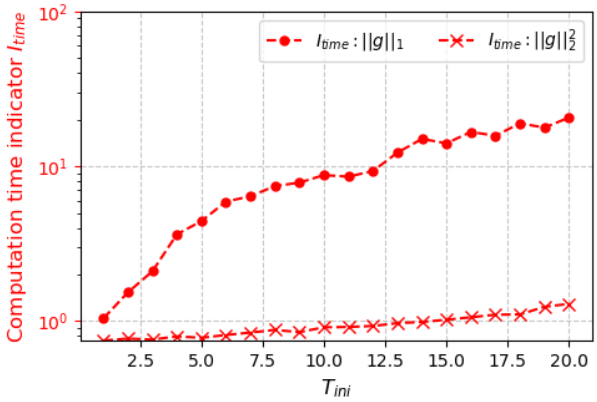

Figure A8: Sensibilité du temps de calcul à l'ordre d'estimation de l'état initial T_{ini}. Les cercles rouges correspondent à la régularisation ℓ_1 et les croix rouges à la régularisation ℓ_2 sur une échelle logarithmique.

Comme illustré sur les Figures A7 et A8, ce compromis se traduit par une courbe en U pour la performance économique et une croissance exponentielle du temps de calcul en fonction de T_{ini}. Il est donc crucial de trouver un équilibre optimal entre précision de prédiction et efficacité computationnelle pour garantir l'applicabilité en temps réel du contrôleur.

Malgré ces difficultés de réglage, les résultats de benchmark démontrent que la DeePC, une fois correctement paramétrée, surpasse à la fois le RC-MPC et la SPC en termes de compromis énergie-confort. Le résultat est illustré plus en détail dans la partie suivante (figure (A10)), où il est comparé aux approches DeePC adaptative.

Extensions adaptatives pour la prise en compte des dynamiques variables

Les modèles thermiques identifiés hors ligne à partir de données historiques peuvent perdre en précision lorsque les conditions opératoires évoluent significativement. Les variations saisonnières des conditions météorologiques, les changements de profils d'occupation, et le vieillissement des équipements HVAC modifient progressivement les paramètres du système tels que les résistances et capacités thermiques, affectant ainsi la qualité des prédictions. Pour maintenir des performances optimales en boucle fermée, des méthodes adaptatives sont nécessaires pour mettre à jour les données utilisées par le contrôleur.

Deux stratégies d'adaptation du DeePC sont explorées dans cette thèse. La première, basée sur une fenêtre glissante, maintient un historique de données de durée fixe qui se déplace au fil du temps en incorporant les nouvelles mesures et en écartant les plus anciennes. Un mécanisme d'évaluation de la qualité des matrices de Hankel est intégré pour détecter les dégradations potentielles de la persistance d'excitation et prévenir ainsi la détérioration des prédictions. La seconde stratégie, fondée sur la réduction récursive, utilise une méthode SVD rapide pour actualiser efficacement la décomposition en valeurs singulières des matrices de Hankel à chaque nouveau pas de temps. Une vérification des

valeurs singulières permet de surveiller la richesse informationnelle des données et de garantir la stabilité numérique du problème d'optimisation.

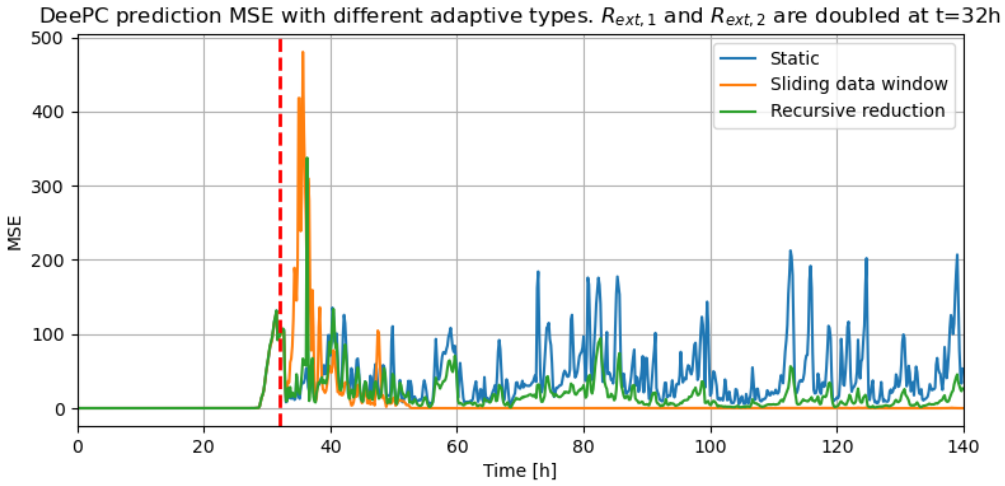

Figure A9: Évolution de l'erreur quadratique moyenne (MSE) pour différents types d'adaptation sur l'exemple académique. Le graphique montre une pièce représentative, car le comportement en cas d'erreur est similaire dans toutes les pièces. Un changement des paramètres thermiques du système intervient à $t = 32$ h (indiqué par une ligne rouge pointillée). L'erreur de prédiction de l'une des pièces est illustrée dans la figure. Bleu : DeePC statique. Orange : DeePC avec fenêtre de données glissante. Vert : DeePC avec réduction récursive. Les courbes montrent la réponse des différentes méthodes d'adaptation face au changement paramétrique dans les deux zones.

Ces approches sont d'abord évaluées sur un exemple académique composé d'un modèle RC à une pièce représentative avec conditions météorologiques constantes (rayonnement solaire et température extérieure fixes). La Figure A9 illustre l'évolution de l'erreur quadratique moyenne (MSE) de prédiction pour trois variantes de DeePC face à un changement soudain des paramètres thermiques du système à $t = 32$ h. Les données d'initialisation des matrices de Hankel et les données de validation sont générées en utilisant un bruit gaussien comme signal d'entrée.

Cette expérience révèle les caractéristiques distinctes des deux stratégies adaptatives. La méthode par fenêtre glissante (courbe orange) génère un pic d'erreur immédiatement après le changement paramétrique, mais parvient à réduire l'erreur à zéro une fois que les matrices de données sont suffisamment actualisées avec les nouvelles mesures. En revanche, la méthode de réduction récursive (courbe verte) offre une transition plus stable et progressive, mais maintient une erreur résiduelle non nulle même en régime établi. La DeePC statique (courbe bleue), quant à elle, accumule progressivement des erreurs sans capacité de correction.

Une étude approfondie du mécanisme d'oubli permet de comprendre l'origine de ces comportements et conduit aux observations suivantes. Premièrement, l'ajout d'un mécanisme d'oubli à la méthode de réduction récursive permet d'atteindre des performances similaires à celles de la fenêtre glissante. L'analyse de différents facteurs d'oubli et tailles de fenêtre glissante révèle un compromis fondamental : l'amplitude du pic d'erreur durant la phase transitoire et le temps nécessaire pour s'adapter aux nouveaux paramètres du système sont inversement corrélés. Deuxièmement, les données collectées durant la phase de transition contiennent intrinsèquement la dynamique du changement paramétrique

lui-même (dans cet exemple, un échelon de la résistance thermique). Par conséquent, les matrices de Hankel durant cette période ne reflètent pas fidèlement les caractéristiques thermodynamiques du système après modification, ce qui engendre le pic d'erreur observé. Pour atténuer cet effet, deux approches sont envisageables : soit mettre en place un mécanisme de sélection pour écarter ces données transitoires, soit choisir une taille de fenêtre suffisamment grande pour diluer leur influence. Dans les deux cas, cela ralentit la vitesse de mise à jour des matrices de Hankel, constituant ainsi le prix à payer pour maintenir la robustesse des prédictions.

De plus, le cas particulier des entrées peu informatives, où le système reçoit des signaux quasi-constants, est également examiné pour évaluer la robustesse des stratégies adaptatives dans ces conditions dégradées. Ces analyses approfondies, incluant l'étude du comportement de l'erreur de prédiction dans différentes phases d'opération, sont détaillées dans le chapitre 4.

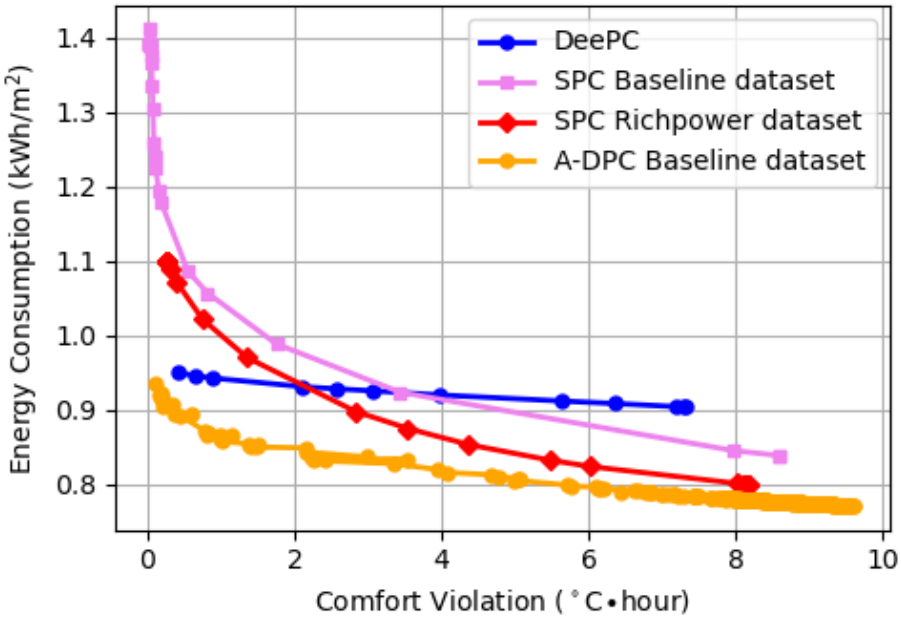

Figure A10: Comparaison des fronts de Pareto des stratégies de contrôle prédictif adaptatives et de référence sur le benchmark MOSA.

Ces solutions sont ensuite évaluées sur le benchmark MOSA (Multi-Occupant Single-zone Air-conditioned). La Figure A10 présente les fronts de Pareto obtenus par différentes stratégies de contrôle, permettant de comparer leurs performances :

- **DeePC avec jeu de données de référence (courbe bleue)** : Offre un bon compromis entre consommation énergétique et confort en exploitant directement les données pour prédire et contrôler, ce qui permet une bonne adaptation aux dynamiques du système sans être limité par les approximations structurelles des modèles paramétriques.

- **SPC avec jeu de données de référence (courbe rose)** : Présente des performances moins favorables en consommation énergétique et en confort, suggérant que le jeu de données de base ne capture pas suffisamment les dynamiques du système.

- **SPC avec jeu de données riche en excitation de puissance (courbe rouge)** : Montre une amélioration significative par rapport au SPC de référence grâce à un jeu de données enrichi, avec une consommation énergétique réduite et un meilleur confort, mais reste moins performant que la DeePC et l'A-DPC.
- **DeePC adaptatif (A-DPC) initialisé par jeu de données de référence (courbe orange)** : Affiche la meilleure performance globale avec la plus faible consommation énergétique et les violations de confort les plus réduites. L'adaptabilité du contrôleur permet de mieux répondre aux variations du système et d'exploiter plus efficacement les apports gratuits en fonction de leur variabilité saisonnière.

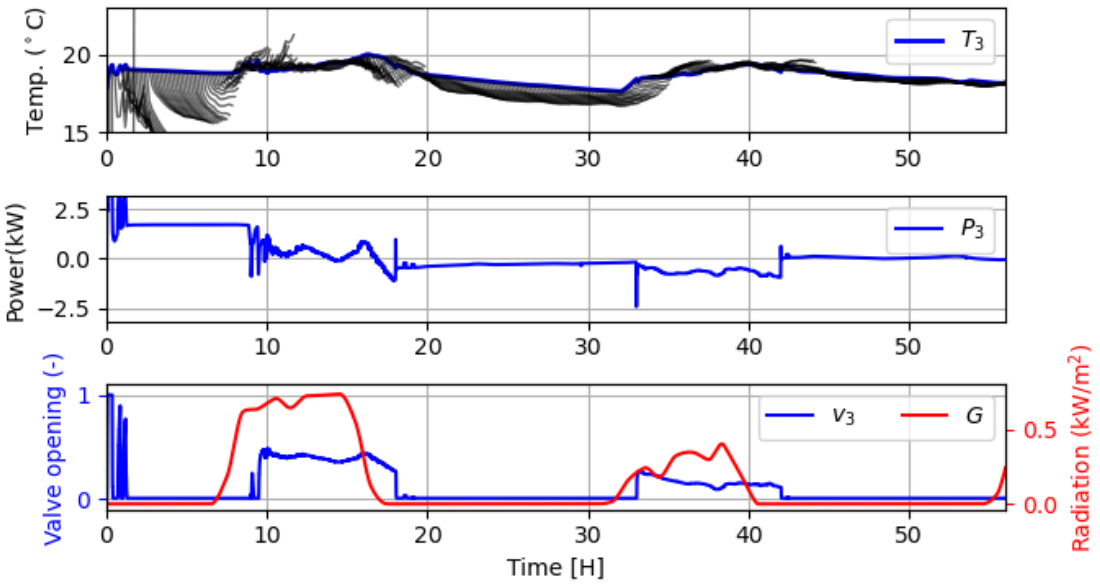

Figure A11: Validation de la qualité de prédiction en boucle fermée sur trois jours d'opération du benchmark MOSA. Le graphique montre une zone représentative. Haut : Trajectoires de la température de zone (en bleu, mesures réelles ; en gris, prédictions). Milieu : Consommation de puissance thermique (en kW), mettant en évidence les motifs diurnes de chauffage. Bas : Ouverture des vannes (en bleu) et perturbation due au rayonnement solaire (en rouge, kW/m^2).

La validation en boucle fermée sur des périodes étendues, illustrée sur la Figure A11, confirme la robustesse des approches adaptatives. Sur une période de trois jours d'opération du benchmark MOSA, en ne considérant que le comportement dynamique d'une zone parmi les 5 disponibles (choisie car représentative), les prédictions de température (trajectoires grises) restent proches des mesures réelles (trajectoires bleues), même durant les périodes d'entrées quasi-constantes où les données apportent peu d'information nouvelle. Les profils de consommation thermique (graphique du milieu) montrent des motifs diurnes typiques, avec des pics de chauffage en début de matinée pour compenser les pertes nocturnes, et des réductions en milieu de journée lorsque les apports solaires sont maximaux. Les actionneurs (vannes de régulation, graphique inférieur) montrent une sollicitation mesurée, révélant une stratégie de contrôle optimisée : le régulateur évite les variations brutales et tire parti de l'inertie thermique du bâtiment pour maintenir un équilibre énergétique.

Synthèse et perspectives

Ce travail de thèse propose un ensemble d'approches de contrôle prédictif basées sur les données pour la gestion thermique des bâtiments, en répondant aux défis posés par l'hétérogénéité des équipements et la disponibilité limitée de données opérationnelles. L'architecture hiérarchique développée permet de découpler la planification thermique de haut niveau, réalisée par la MPC superviseur, de la régulation des équipements de bas niveau, assurée par des contrôleurs PI ajustés par VRFT. Cette séparation améliore la robustesse et la modularité du système de contrôle, facilitant ainsi son déploiement sur différents types de bâtiments.

Trois approches de contrôle prédictif sont comparées: le RC-MPC, qui exploite des modèles physiques informés par l'énergie et offre des performances robustes avec un effort de réglage minimal; la SPC, qui repose sur des modèles boîte noire identifiés par sous-espaces et offre un compromis intéressant entre besoins en données et complexité algorithmique; et la DeePC, qui opère directement sur les données sans identification de modèle explicite, atteignant les meilleures performances mais au prix d'un effort de réglage non-trivial et de ressources de calcul accrues. Le Tableau A12 synthétise les caractéristiques de ces trois approches selon plusieurs critères clés.

Figure A12: Évaluation comparative des approches de contrôle prédictif basées sur les données pour la gestion thermique des bâtiments

Critère	**RC-MPC**	**SPC**	**DeePC**
Effort de modélisation	Élevé	Bas	Bas
Besoin en données	Moyen	Moyen	Élevé
Complexité algorithmique	Moyen	Moyen	Élevé
Temps d'optimisation	Bas	Moyen	Moyen
Mémoire de calcul	Bas	Bas	Élevé
Effort de réglage	Bas	Moyen	Élevé
Performance atteignable	Élevée	Moyen–Élevée	Élevée

Le RC-MPC est recommandé quand l'expertise en modélisation physique et la connaissance du domaine sont disponibles. La SPC convient quand les ressources de données et de calcul sont modérés. La DeePC est pertinente quand des données abondantes et une capacité de calcul élevée sont accessibles.

Les extensions adaptatives développées pour prendre en compte les dynamiques variables du système représentent une contribution significative de cette thèse. Les résultats démontrent que l'adaptation en ligne des modèles ou des données permet d'améliorer substantiellement les compromis énergie-confort sous conditions opératoires changeantes, ce qui est essentiel pour garantir des performances optimales tout au long de l'année.

Les perspectives de ce travail incluent l'extension des méthodes proposées à des systèmes de contrôle multi-bâtiments pour la gestion de districts énergétiques, l'intégration de prévisions probabilistes pour mieux gérer les incertitudes sur les prévisions météorologiques et d'occupation, et l'incorporation de contraintes de réseau électrique pour faciliter la participation des bâtiments aux marchés de flexibilité énergétique. De plus, le développement d'outils logiciels open-source permettrait de faciliter l'adoption de ces technologies par l'industrie et de contribuer à la transition énergétique du secteur du bâtiment.

Articles de revues

- Yuqi Liu, Pauline Kergus, Fabien Claveau, Philippe Chevrel, and Bruno Lacarrière. Building thermal control: Hierarchical design from limited data using gray-box or black-box internal models for model predictive control. *Journal of Building Engineering*, 110:112936, 2025.

Communications avec actes

- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Hierarchical Model Predictive Control for Building Thermal Management: A mixed model-based and data-based approach. In *Proc. 11th International Conference on Control, Decision and Information Technologies (CoDIT)*, Split, Croatia, 2025. IEEE.
- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Hierarchical Data-enabled Predictive Control for Building Thermal Management. *IFAC Control Application of Optimization Workshop*, 2025.

Revues en préparation

- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Online Adaptation of Data-Enabled Predictive Control: Pro and Cons for Optimized Building HVAC Management. En préparation.

Software implementation

- Yuqi LIU, HierCon: A Python Toolkit for Hierarchical and Optimized Thermal Management in Buildings. Internal Report, IMT Atlantique, 2025.

Table of Contents

Chapter I: Introduction **28**

1 Challenges in Building Thermal Management 28
- 1.1 Equipment Diversity and Control Challenges 29
- 1.2 Limited Data Availability in Building Operations 30
- 1.3 Methodology Overview and Contributions 30
- 1.4 Thesis Organization 33

2 Building Thermal Control Validation Framework 35
- 2.1 BOPTEST Framework Architecture and Implementation 35
- 2.2 Building Thermal Dynamics and Physical Modeling 36
- 2.3 Multizone Test Cases: MZRH and MOSA 39
- 2.4 Multi-objective Formulation and Performance Metrics 42

3 Hierarchical Control Methodology 43
- 3.1 Hierarchical Control Architecture 44
- 3.2 Implementation Workflow and Design Procedure 45

Chapter II: Hierarchical MPC with Gray-box and Black-box Models **48**

1 Literature Review on MPC for Buildings 48
- 1.1 MPC Classification and Implementation Approaches 48
- 1.2 Gray-box model vs. Black box model for buildings 50

2 System Identification for Prediction Models 54
- 2.1 Gray-box Resistance-Capacitance Model Identification 54
- 2.2 Practical Implementation Considerations 61

3 Hierarchical MPC Design 62
- 3.1 Supervisory-level MPC Formulation 63
- 3.2 Tracking-level Controller Tuning via VRFT 66

4 Simulation Validation and Performance Evaluation 69
- 4.1 Open-loop Validation of Indoor Climate Models 69
- 4.2 Tracking Controller Performance Assessment 71
- 4.3 Closed-loop Hierarchical Control Performance 72

Chapter III: Model-free Predictive Control with DeePC **79**

1 Preliminaries on Data-driven Control 79
- 1.1 Direct Data-driven Control Framework 80
- 1.2 Data-enabled Predictive Control 83

1.3 Subspace Predictive Control . . . 87
1.4 Comparison of DeePC and SPC . . . 90
2 DeePC Implementation and MOSA Validation . . . 92
2.1 Hierarchical DeePC Implementation Setup . . . 93
2.2 Closed-loop Performance and Comparative Analysis . . . 95
2.3 Sensitivity Analysis of DeePC Hyperparameters . . . 99
Chapter IV: Adaptive Data-driven Predictive Control . . . 105
1 Preliminaries on Adaptive Predictive Control . . . 105
1.1 Motivation and Research question . . . 105
1.2 Literature review on Adaptive MPC . . . 106
2 Adaptive DPC: Classification and Preliminaries . . . 110
2.1 Adaptive Data Update Strategies . . . 110
2.2 Data Quality Assessment . . . 114
3 Adaptive Strategy Validation . . . 117
3.1 Numerical Case Study: RC Parameter Variance . . . 117
3.2 MOSA benchmark validation . . . 128
Chapter V: Thesis Conclusions and Perspectives . . . 132
1 Synthesis and Contributions . . . 132
1.1 Hierarchical Framework and Validation Protocol . . . 132
1.2 Data-Driven Predictive Control Approaches . . . 133
1.3 Adaptive Mechanisms . . . 135
1.4 Thesis Conclusion . . . 136
1.5 Valorisation and Scientific Contributions . . . 136
2 Perspectives . . . 137
2.1 Perspectives on Data-Driven Predictive Control . . . 137
2.2 Perspectives on Building Control Implementation . . . 138
Appendix A: Extension to 2nd order RC model . . . 140
Appendix B: SPC and DDPC comparison for LPV systems . . . 145

List of Figures

I.1 The BOPTEST concept [1]. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. . . . 36
I.2 An example representation of a room with a RC network. Source: [2], with permission from Elsevier, as retained by the author). . . . 36
I.3 MZRH test case layout and hydronic circuit . . . 40
I.4 MZRH test case radiation and exterior temperature profile. . . . 41
I.5 MOSA test case benchmark . . . 42
I.6 Hierarchical control scheme architecture . . . 44
I.7 Implementation workflow . . . 45
II.1 Architecture paradigms of MPC solution for buildings . . . 49
II.2 1st order RC circuit multi-zone model . . . 56
II.3 Multi-zone 1st order RC model. . . . 57
II.4 The proposed hierarchical control scheme applied to MZRH test case. . . . 62
II.5 Thermal power response characteristics . . . 64
II.6 VRFT process block diagram . . . 67
II.7 Indoor climate model prediction validation . . . 70
II.8 Error propagation with growth of prediction horizon . . . 71
II.9 Historical dataset length dependency . . . 71
II.10 VRFT controller validation . . . 72
II.11 1RCMZ vs 4SID performance comparison . . . 73
II.12 Zoom-in comparison of MPC formulations . . . 74
II.13 MPC Prediction quality comparison . . . 75
II.14 Radiator power tracking performance in MZRH case . . . 75
II.15 KPI map for control schemes in MZRH case . . . 75
III.1 Venn diagram of data-driven predictive control approaches . . . 80
III.2 Hierarchical DeePC control scheme validation results . . . 96
III.3 Pareto fronts comparing DeePC ($T_{ini} = 5$) and RC-MPC control schemes . . . 97
III.4 Pareto fronts comparing DeePC and SPC control schemes with different dataset. Both method initial condition estimation order is 5. . . . 97
III.5 Economic performance vs initial state order evaluation in MOSA case . . . 99
III.6 Computation time vs initial state order in MOSA case . . . 100
III.7 Regularization weights sensitivity analysis in MOSA case . . . 101
IV.1 Second-order RC network topology . . . 119
IV.2 Initial dataset for adaptive DeePC predictor . . . 120
IV.3 Data-driven prediction validation . . . 120

IV.4 Static DeePC under parameter shift . . . 121
IV.5 MSE evolution for non-adaptive prediction . . . 122
IV.6 MSE comparison for adaptive strategies . . . 122
IV.7 Forgetting factor sensitivity analysis . . . 124
IV.8 Validation inputs with excitation degradation . . . 125
IV.9 MSE under excitation degradation . . . 126
IV.10 Performance with data quality indicators . . . 127
IV.11 Closed-loop prediction quality . . . 129
IV.12 Pareto front comparison adaptive control . . . 130
A.1 Representation of a zone in 2nd order single zone RC model. . . . 140

List of Tables

I.1 Overview of the proposed methodology structure and investigated variants. . . . 31
II.1 VRFT tuned PI controller parameters . . . 71
II.2 Rate-constrained MPC parameters used in test case . . . 72
III.1 Fixed experiment configuration in case study . . . 94
III.2 PI parameters for each zone tuned by VRFT . . . 95
IV.1 Interpretation of data quality indicators: high vs low values . . . 118
IV.2 Adaptive mechanism embedded in figure IV.6 experiments . . . 122
IV.3 Adaptive mechanism embedded in figure IV.7 experiments . . . 125
IV.4 Adaptive mechanism embedded in figure IV.9 experiments . . . 126
IV.5 Adaptive mechanism embedded in figure IV.9 experiments . . . 127
IV.6 Adaptive DeePC controllers configuration for MOSA benchmark. Adaptive strategy is algorithm 1. DeePC and simulation configuration follows table III.1 . . . 128
V.1 Comparative evaluation of data-driven predictive control approaches for building thermal management. . . . 134

Acronyms

Acronym	Full Term
4SID	Subspace State Space Identification
ADeePC	Adaptive Data-enabled Predictive Control
AHU	Air Handling Unit
API	Application Programming Interface
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers
ASPC	Adaptive Subspace Predictive Control
BOPTEST	Building Optimization Testing Framework
DeePC	Data-enabled Predictive Control
EPBD	Energy Performance of Buildings Directive
FMI	Functional Mock-up Interface
FMU	Functional Mock-up Unit
HVAC	Heating, Ventilation and Air Conditioning
I/O	Input/Output
IBPSA	International Building Performance Simulation Association
LMTD	Log Mean Temperature Difference
MOSA	Multi-zone Office Simple Air
MPC	Model Predictive Control
MZRH	Multi-zone Residential Hydraulic
PID	Proportional-Integral-Derivative
PRBS	Pseudo-Random Binary Sequence
RBC	Rule-Based Control
RC	Resistance-Capacitance
REST	Representational State Transfer
SPC	Subspace Predictive Control
SVD	Singular Value Decomposition
VAV	Variable Air Volume
VRFT	Virtual Reference Feedback Tuning
ZOH	Zero-order Hold

Mathematical Nomenclature

Symbol	Description
State and Control Variables	
v	Valve opening percentage (control input)
v_i	Valve opening for zone i
P	Thermal power vector [W]
$P_i(t)$	Thermal power delivered to zone i [W]
T	Temperature vector [K]
$T_i(t)$	Air temperature in zone i [K]
T_{ext}	Exterior air temperature [K]
d	Weather disturbance vector
G	Solar radiation [W/m^2]
u	General control input vector
y	General output/measurement vector
x	State vector
System Parameters	
n_z	Number of zones
$N_{\text{ad},i}$	Set of adjacent zones to zone i
C	Thermal capacitance [J/K]
R	Thermal resistance [K/W]
$C_{z,i}$	Thermal capacitance of zone i
$R_{i,j}$	Thermal resistance between zones i and j
$R_{i,\text{ex}}$	Thermal resistance to exterior for zone i
K_P, K_I, K_D	PID controller gains (proportional, integral, derivative)
θ_{PID}	PID parameter vector
Time and Horizons	
t	Continuous time [s]
k	Discrete time step
t_0, t_f	Initial and final time
T_s	Sampling time
T_{PID}	PID discrete time step
N	Prediction horizon
T_{ini}	Initial data length/horizon
ℓ	Data window length

Symbol	Description
Optimization and Cost	
J_{energy}	Energy consumption objective
$J_{\text{discomfort}}$	Thermal discomfort objective
J_{total}	Total cost function
λ	Weighting factor for discomfort penalty
$\underline{T}(t)$	Lower comfort temperature boundary [K]
Matrix and Data Notation	
H	Hankel matrix
H_k	Hankel matrix at time step k
U_f	Future input Hankel matrix
U_p	Past input Hankel matrix
Y_f	Future output Hankel matrix
Y_p	Past output Hankel matrix
$\Pi_{U_f}^{\perp}$	Projection onto orthogonal complement of U_f
I	Identity matrix
$(\cdot)^{\dagger}$	Moore-Penrose pseudoinverse
$(\cdot)^{\top}$	Matrix transpose
System Identification	
A, B, C, D	State-space matrices
a_{ij}	Element (i, j) of matrix A
n	Model order / number of states
$T_{i,j}^{m}$	Temperature of adjacent zone j to zone i, m-th neighbor
Data and Sequences	
u_d	Historical input data sequence
y_d	Historical output data sequence
u_d^k	Input data up to time k
y_d^k	Output data up to time k
w	Vector in kernel space
DeePC and Regularization	
g	Trajectory combination vector
σ_y	Slack variable for output matching
λ_y	Regularization weight on slack variable
λ_g	Regularization weight on trajectory vector g
T_d	Dataset length (number of data samples)
h	Trajectory length in Hankel matrix
p	Norm order ($p \in \{1, 2\}$)

Symbol	Description
Adaptive Control	
γ	Forgetting factor
N_{eff}	Effective memory length
I_{match}	Prediction error indicator
I_{info}	Data informativity indicator
I_{oper}	Operational relevance indicator
c_k	New data column candidate at time k
ϵ_{tol}	Tolerance threshold for numerical rank
ϵ_{info}	Informativity threshold
α_{off}	Off-service data weight
$\mathcal{T}_{\text{service}}$	Set of service hours
N_w	Sliding window length for indicator averaging
Other	
$\mathcal{N}$	Null space / kernel
$\text{span}(\cdot)$	Span of a set of vectors
$\|\cdot\|$	Norm operator
$\max(\cdot)$	Maximum function
$\emptyset$	Empty set

Chapter I

Introduction

1 Challenges in Building Thermal Management 28
1.1 Equipment Diversity and Control Challenges 29
1.2 Limited Data Availability in Building Operations 30
1.3 Methodology Overview and Contributions 30
1.4 Thesis Organization 33
2 Building Thermal Control Validation Framework 35
2.1 BOPTEST Framework Architecture and Implementation 35
2.2 Building Thermal Dynamics and Physical Modeling 36
2.3 Multizone Test Cases: MZRH and MOSA 39
2.4 Multi-objective Formulation and Performance Metrics 42
3 Hierarchical Control Methodology 43
3.1 Hierarchical Control Architecture 44
3.2 Implementation Workflow and Design Procedure 45

1 Challenges in Building Thermal Management

Buildings account for 30-40% of global energy consumption [3]. Thermal management systems—including Heating, Ventilation and Air Conditioning (HVAC)—represent approximately one-third of this use [4]. The Energy Performance of Buildings Directive [5] identifies thermal management efficiency as a key performance indicator. This calls for advanced control strategies balancing occupant comfort with energy efficiency.

Model Predictive Control (MPC) has shown considerable potential for building thermal management [6]. By predicting a building's indoor climate evolution, MPC exploits thermal inertia to improve system performance. Studies demonstrate that MPC outperforms classic rule-based control in multi-objective planning [7], achieving energy consumption reductions while maintaining thermal comfort [8, 9, 10, 11].

The performance of MPC depends on the accuracy of its prediction models. First-principles modeling approaches derive system dynamics from physical laws and conservation equations. Notable contributions [12, 13, 14] developed detailed simulation models using Modelica for MPC applications. Alternatively, other investigations [15, 16, 17, 18, 19] pursued data-based modeling to reduce modeling effort. Despite these advances, modeling remains a critical bottleneck in MPC deployment. The specialized knowledge, resources, and time required for developing and validating models diminish MPC's applicative advantages [20]. This bottleneck has motivated extensive research into data-driven control methods that learn system dynamics directly from measured data [16, 17, 18, 19, 21, 22]. By eliminating the need for first-principles modeling, these approaches reduce engineering effort and

enable faster deployment across diverse building configurations. While data-driven methods reduce modeling effort, two key challenges limit their practical deployment in building thermal management:

1. **Equipment Diversity**: HVAC systems vary widely in design, behavior, and control requirements. Control methods trained on one system type may not generalize well to different buildings with distinct equipment configurations. This specificity creates difficulties in developing solutions that work across different buildings without extensive adaptation.

2. **Limited Data Availability**: Data-driven methods typically require informative training data to achieve reliable performance. Real buildings provide only historical operational data, constrained by comfort requirements and normal operation. This data may not cover all conditions that the control system needs to explore for optimization. Models trained on limited operational data alone may not perform reliably when the system operates beyond typical ranges.

These challenges call for control methods that reduce modeling complexity while maintaining reliable performance under realistic data conditions. The following sections examine how equipment diversity and data limitations manifest in building thermal management.

1.1 Equipment Diversity and Control Challenges

HVAC systems encompass diverse solutions to maintain indoor thermal conditions and air quality. In large-scale buildings, this often involves Air Handling Unit and Variable Air Volume systems for centralized climate control. For residential and smaller commercial applications, conventional solutions include hydraulic circulation with radiators, heat pumps, and thermal energy storage systems. The diversity of HVAC equipment and their control requirements presents a significant challenge in developing unified management strategies. This challenge stems from the building-specific nature of HVAC installations, where system scale, network topology, and equipment configurations vary significantly across different buildings. Moreover, these systems typically exhibit nonlinear behaviors and are subject to operational constraints that further complicate their control.

Addressing this equipment diversity challenge, previous research has explored various thermal control methodologies. For instance, [23] built an environment with indoor climate and HVAC models for hierarchical reinforcement learning. [24] used data generated from physical models to train data-based models for adaptive MPC. Although these approaches provide valuable insights, they typically result in system-specific solutions. This specificity necessitates the development of unique device models for each HVAC configuration, which remains a barrier to widespread implementation.

In contrast to the methods mentioned above, which seek to model the thermal inertia of buildings with HVAC actuators, a number of research assume an ideal response of HVAC to avoid engaging in its modeling [25, 26, 27, 28, 29]. These works reformulate the control problem to operate directly with heating and cooling power as control variables, abstracting away the underlying actuation mechanisms such as valves, dampers in system like radiators, air conditioners. These approaches maintains independence from the particular type of HVAC system installed, offering a more generalized solution. While this strategy significantly reduces the complexity of model development for predictive control design, it introduces potential challenges in practical implementation, particularly in translating the

computed power signals into actual equipment control commands that achieve the desired thermal output.

1.2 Limited Data Availability in Building Operations

Data quality and excitation play a critical role in building control design. Open-loop experiments with persistently exciting inputs, such as pseudo-random binary sequences or white/colored noise, are known to provide reliable data-driven models [30]. However, such experiments are often impractical in real building thermal systems. Instead, historical operational and maintenance data, though limited in amplitude and frequency by closed-loop operation, are more realistically available. Under these constraints, the choice between gray-box and black-box modeling becomes essential, as each relies differently on data quality, quantity, and characteristics.

Gray-box modeling reduces reliance on large amounts of data by informing prior knowledge of the model structure [6, 31, 32]. It usually entails a vast simplification and linearization of system physics. Equivalent Resistance - Capacitance (RC) model is a typical gray-box model for buildings thermal behavior. It translates building thermodynamic to a given electric circuit topology. Its capability for modeling indoor climate is extensively studied [33, 34, 35, 36].

Black-box models learn system dynamics directly from measured data without physical assumptions. Linear models, such as autoregressive-based models and 4SID models, offer computational efficiency and interpretability [16, 17, 37]. A comparison of these models is provided in [17], showing that 4SID offers better approximation performance and requires less modeling effort than autoregressive-based models. Nonlinear models such as Artificial Neural Networks and Gaussian Processes often outperform linear models, providing flexible input-output mapping and uncertainty quantification. However, they require substantial training data to achieve reliable performance [18, 19]. When training data are limited in amplitude and frequency, black-box models may struggle to extrapolate beyond observed conditions [38]. This limitation becomes particularly critical when models are embedded in MPC, where the optimization may explore operating regions not well represented in the training data.

Efficient predictive models require that data collection conditions closely match deployment conditions [39]. This includes sampling time, prediction horizon, and operating points. Historical building data typically reflect occupant comfort constraints, whereas MPC optimization may explore a broader range of scenarios to identify the optimum. As a result, predictive models trained solely on operational data may be insufficient. It is therefore essential to validate the proposed control strategies using only historical operational data, in order to demonstrate the robustness of the methodology under the practical conditions considered.

1.3 Methodology Overview and Contributions

The objective of this thesis is to develop control solutions that address the challenges above. This work explores how data-driven predictive control methods can be applied to building thermal management. By leveraging established methodologies and evaluating their performance under realistic operational constraints, this research aims to bridge the gap between model-based/model-free predictive control, and practical deployment in building HVAC systems. The contribution lies in extracting insights from comparative analysis, identifying the strengths and limitations of different

approaches, and providing guidance for practitioners on deploying appropriate methods for real-world application. The challenges of equipment diversity and limited data availability motivate the combination of specific control approaches. This work studies existing data-driven methods that are potentially capable of handling these constraints.

This work adopts a hierarchical control architecture that addresses these challenges at different levels. The architecture separates thermal dynamics management at the supervisory level from equipment control at the tracking level, allowing each level to employ methods best suited to its specific requirements. The proposed hierarchical control framework contains multiple components and variants investigated throughout this thesis. Table I.1 presents an overview of the key component menu.

Supervisory level	
Model predictive control	Direct Data-driven Predictive Control
• Gray-box model: RC • Black-box model: 4SID	• Data Enabled Predictive Control
Adaptive Mechanism	
• Sliding data window • Recursive reduction	
Tracking level	
• Virtual Reference Feedback Tuning	
Benchmark	
• Office: Air Conditioner • Residence: Hydronic Radiator	

Table I.1: Overview of the proposed methodology structure and investigated variants.

At the supervisory level, three predictive control approaches are investigated to manage thermal dynamics under limited data availability:

- **Gray-box MPC using RC models**: Incorporates physical structure into data-driven models to constrain the optimization problem and provide robustness for model extrapolation beyond training data. The RC model encodes thermal principles while learning parameters from data. This combination reduces data requirements compared to pure black-box methods while maintaining interpretability and reliability.

- **Black-box MPC using 4SID**: Subspace identification methods learn system dynamics entirely from data without physical assumptions. 4SID offers computational efficiency and avoids the modeling effort required for first-principles approaches, while still providing state-space representations suitable for predictive control.

- **Model-free DeePC**: Data-enabled Predictive Control offers an alternative paradigm that avoids explicit modeling. DeePC uses input-output data directly for prediction and control, eliminating the modeling bottleneck while preserving the predictive control framework. This approach is particularly attractive when system complexity makes even data-driven modeling impractical.

This comparative study evaluates how different modeling paradigms perform under limited data conditions using only historical operational data. The RC model incorporates thermal physics to

improve extrapolation, while 4SID and DeePC rely purely on data patterns. By comparing these approaches, this work provides practical insights into their respective advantages and limitations for building thermal management under realistic data constraints.

At the tracking level, VRFT [40] is used to address the equipment diversity challenge. VRFT tunes controller parameters using only input-output data to match a desired closed-loop reference model, without requiring detailed HVAC system models. This model-free characteristic makes VRFT particularly suitable for handling diverse HVAC equipment like radiators, heat pumps, and air conditioners. VRFT is applied to tune PID controllers, which are well-understood, easy to analyze, and widely accepted in industrial practice. VRFT has demonstrated its precision and scalability in handling complex thermodynamic systems within temperature uniformity control systems [41]. However, its applicability to building thermodynamic systems has not been sufficiently explored. This work incorporates it into the tracking layer methodology as a design approach for managing equipment diversity.

The hierarchical separation through power signals has been successfully demonstrated in previous studies [42, 33, 43]. By abstracting the supervisory control to operate on thermal power rather than actuator positions, the framework maintains independence from specific HVAC equipment types. This power-based abstraction addresses the equipment diversity challenge by decoupling high-level thermal planning from low-level actuation details.

Furthermore, the control schemes investigated in this work are extended to adaptive versions through online parameter updates. This extension enables the controllers to adjust to time-varying environmental factors such as seasonal weather changes and solar radiation variations. The adaptive mechanisms are particularly important for long-term operation where system characteristics evolve due to seasonal variations, occupancy patterns, and equipment aging.

This framework balances multiple practical considerations: modeling complexity is reduced through data-driven approaches, equipment diversity is handled through hierarchy separation and model-free tuning, data limitations are addressed through physics-informed models and careful validation on operational data, and time-varying dynamics are managed by adaptive mechanisms. The following sections present detailed reviews of the selected control methods, examining their theoretical foundations, practical implementations, and applicability to building thermal management.

The methodology presented in this thesis makes several contributions to building thermal management:

- A modular hierarchical control framework that addresses equipment diversity and data limitation challenges through power-signal separation, enabling equipment-agnostic supervisory control.
- Integration of data-driven techniques at both control levels, introducing VRFT to building thermal management for model-free controller tuning at the tracking level.
- Rate-constrained predictive control formulation that accounts for asymmetric power dynamics in high-thermal-inertia HVAC equipment, with constraints identified directly from operational data.
- Comparison of gray-box RC and black-box 4SID models through both open-loop prediction validation and closed-loop control performance evaluation under realistic data conditions.
- Benchmarking of DeePC DeePCagainst model-based MPC, demonstrating superior energy efficiency and providing practical hyperparameter guidelines for deployment.

- Adaptive mechanisms using sliding data window and SVD-based recursive reduction for Hankel matrix updating in DeePC and Subspace Predictive Control, enabling online adaptation to time-varying conditions.

Journal Articles

- Yuqi Liu, Pauline Kergus, Fabien Claveau, Philippe Chevrel, and Bruno Lacarrière. Building thermal control: Hierarchical design from limited data using gray-box or black-box internal models for model predictive control. *Journal of Building Engineering*, 110:112936, 2025.

Conference Proceedings

- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Hierarchical Model Predictive Control for Building Thermal Management: A mixed model-based and data-based approach. In *Proc. 11th International Conference on Control, Decision and Information Technologies (CoDIT)*, Split, Croatia, 2025. IEEE.
- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Hierarchical Data-enabled Predictive Control for Building Thermal Management. *IFAC Control Application of Optimization Workshop*, 2025.

Journals in Preparation

- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. DeePC and SPC for Building Energy Management: A Comparative Evaluation of Their Suitability for System Evolving Dynamic. In preparation.
- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Online Adaptation of Data-Enabled Predictive Control: Pro and Cons for Optimized Building HVAC Management. In preparation.

1.4 Thesis Organization

This manuscript presents the Ph.D. research in five chapters. The development and validation of data-driven control solutions for building thermal management is presented in the following order.

Chapter I: Introduction This chapter provides an overview of the building energy control domain and establishes the research context. A literature review identifies the key challenges in data-driven control: equipment diversity and limited data availability. Based on this analysis, the research objectives are formulated, aiming to explore and validate data-driven control methods for practical building applications. The proposed hierarchical control framework consists of two levels: a supervisory-level predictive control for thermal dynamics management and a tracking-level control for equipment regulation. The validation methodology uses the Building Optimization Testing Framework

(BOPTEST) as the simulation benchmark. Two test cases from BOPTEST are used throughout the thesis to evaluate the proposed approaches. The control solutions are integrated into the 'hiercon' toolbox, a Python-based framework developed in this thesis for co-simulating with benchmark. The general methodology, workflow, and framework structure are presented in this chapter, while detailed mathematical formulations and validation results are distributed across the subsequent chapters.

Chapter II: Hierarchical MPC with Gray-box and Black-box Models This chapter investigates predictive control at the supervisory level, comparing physics-informed and purely data-driven modeling approaches. Two system identification methods are examined: 1) the gray-box Resistance-Capacitance model that incorporates thermal physics, 2) the black-box 4SID method, which learns dynamics entirely from data. Both models are trained using limited historical operational data and integrated into a MPC framework with rate constraints to handle actuator dynamics. At the tracking level, VRFT is applied to tune decentralized PID controllers without requiring explicit HVAC system models. The chapter presents detailed mathematical formulations for system identification, MPC design, and VRFT-based controller tuning. Validation on both MZRH and MOSA test case demonstrates that while both modeling approaches achieve similar open-loop prediction accuracy, their closed-loop performance differs significantly. The RC model achieves 3-12% energy savings compared to the 4SID model while maintaining comparable thermal comfort. The gray-box model demonstrates superior closed-loop performance under conditions of limited data quality.

Chapter III: Model-free Predictive Control with DeePC This chapter extends the investigation to model-free predictive control using DeePC. Unlike the indirect data-driven approach in chapter II, where models are first identified and then used for prediction, DeePC directly uses input-output data for trajectory prediction based on behavioral systems theory [44]. The chapter presents the theoretical foundations of DeePC [45], including Willems' fundamental lemma and the construction of Hankel matrix from persistently exciting data. The hierarchical framework is adapted to integrate DeePC at the supervisory level while maintaining VRFT-tuned PID controllers at the tracking level. Detailed hyper-parameter analysis is provided to guide practical implementation, examining the effects of data horizon length, regularization parameters, and slack variable weights on control performance. Validation on the MOSA handling test case demonstrates that DeePC achieves superior energy efficiency and comfort management compared to conventional RC-based MPC, while completely eliminating the modeling process. The chapter concludes with practical guidelines for DeePC deployment in building control applications.

Chapter IV: Adaptive Data-driven Predictive Control This chapter addresses time-varying system dynamics through adaptive control mechanisms. Building thermal systems exhibit evolving characteristics due to seasonal weather patterns, occupancy variations, and equipment aging. The control schemes developed in Chapters 2 and 3 may suffer from performance degradation during long-term operation when system characteristics drift from their initial training conditions. This chapter develops adaptive extensions of both model-based and model-free predictive control approaches. Adaptive Subspace Predictive Control (ASPC) continuously updates the 4SID model through recursive identification using sliding data windows. Adaptive DeePC (ADeePC) directly refreshes the Hankel data matrix without explicit model identification. The chapter presents different data update strategies and evaluates their computational efficiency and adaptation performance. Theoretical analysis establishes the relationship between ASPC and ADeePC, showing their duality through the shared reliance on

data matrix column spaces. MOSA demonstrates that adaptive schemes maintain control performance under seasonal variations and long-term system drift, providing robustness that static approaches cannot achieve.

Chapter 5: Conclusions and Perspectives This chapter synthesizes the key findings from the studies above and discusses their value for practical building thermal management. The main contributions are summarized: evaluation of data-driven control approaches under realistic constraints, practical guidelines for application. Conduct an analysis of the trade-offs among modeling complexity, data requirements, and control performance. The chapter also discusses the limitations of current approaches and proposes future research directions, including the integration of uncertainty quantification, hyperparameter tuning, and real-world experimental validation. Final remarks shows the practical value of bridging advanced control theories with building energy applications through exploration and comparative analysis.

2 Building Thermal Control Validation Framework

2.1 BOPTEST Framework Architecture and Implementation

Numerous test cases are available in the literature on thermal control in buildings. Real building experiments constrain methodological research due to several limitations: they are time-consuming, expensive, subject to occupant comfort requirements, and difficult to reproduce under identical conditions. The Building Optimization Testing Framework (BOPTEST) [1] addresses these challenges by providing a widely adopted standardized platform for evaluating building control strategies. Developed through international collaboration by International Building Performance Simulation Association (IBPSA) Project 1 researchers, BOPTEST offers realistic test cases spanning diverse building types, including commercial and residential buildings, single-zone and multi-zone configurations, and various HVAC systems such as hydronic and air-based solutions. This framework enables rigorous, reproducible benchmarking of control methodologies under realistic operating conditions, making it well-suited for systematic validation and performance comparison.

BOPTEST employs a modular architecture that separates building simulation from control algorithm implementation, facilitating controller evaluation. The framework is constructed using Modelica, an equation-based, object-oriented modeling language particularly suited for multi-domain physical systems. Building models are developed using established Modelica libraries, including IDEA [46] for district energy systems and the Buildings library [47] for building energy systems. These libraries provide validated component models for building envelopes, HVAC equipment, and control systems that have been extensively tested against experimental data.

As shown in Figure (I.1), BOPTEST employs an architecture where users can test their control algorithms on chosen building emulators. The simulation models are compiled into Functional Mock-up Units following the FMI standard, enabling platform-independent co-simulation. The building model executes independently while exchanging data with external control algorithms through a RESTful HTTP API deployed in Docker containers. This architecture allows control algorithms developed in Python to request measurements, send control actions, and advance simulation time.

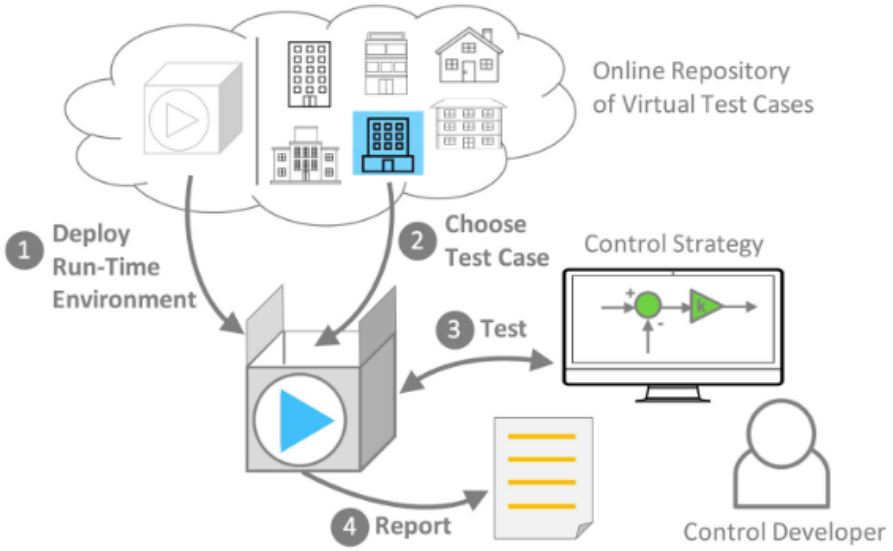

Figure I.1: The BOPTEST concept [1]. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Upon completion, the framework generates a report containing performance metrics evaluated based on the emulator's operation.

2.2 Building Thermal Dynamics and Physical Modeling

BOPTEST test cases capture the essential complexities of real building thermal systems through high-fidelity physical models. The modeling approach incorporates multiple nonlinear phenomena and coupled multi-domain dynamics that challenge control system design.

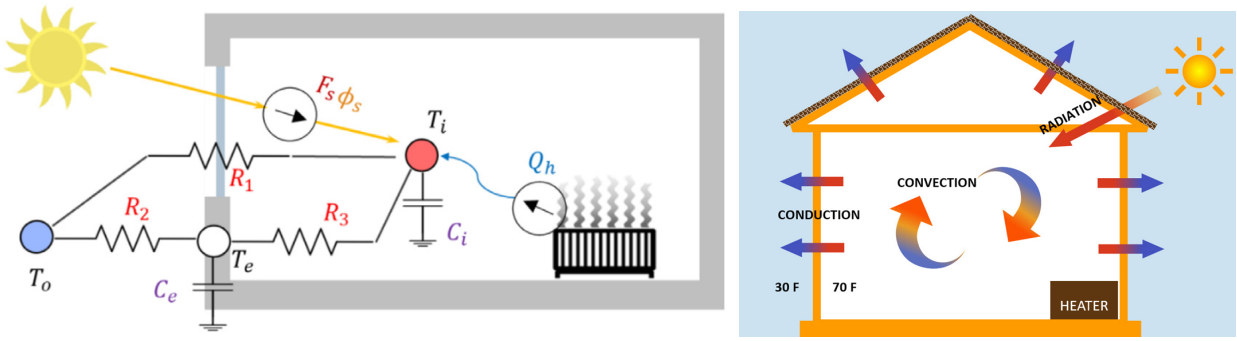

Figure I.2: An example representation of a room with a RC network. Source: [2], with permission from Elsevier, as retained by the author).

Figure I.2 illustrates thermal dynamics in building indoor climate through a resistance-capacitance (RC) network analogy. The outdoor temperature T_o connects to the indoor temperature T_i through a series of thermal resistances (R_1, R_2, R_3 representing building envelope layers), with thermal

capacitances (C_e, C_i) capturing heat storage in the envelope and indoor air respectively. Solar radiation $F_s\phi_s$ enters through the window as a heat source, while heating input Q_h from the HVAC system provides controllable thermal energy to maintain indoor comfort. The envelope temperature node T_e represents an intermediate thermal mass between exterior and interior conditions. This representation captures the essential physics governing building thermal response: conductive heat transfer through multi-layer walls, thermal energy storage in building mass and air, solar heat gains, and HVAC heating, which together determine the indoor temperature dynamics.

Conductive and convective heat transfer. Conductive heat transfer through multi-layer walls, floors, and roofs uses lumped parameter resistance-capacitance networks that capture thermal mass effects and spatial temperature gradients, enabling accurate representation of transient heat storage and release. Convective heat transfer employs empirical correlations that depend nonlinearly on surface temperatures, air temperatures, and velocities. Interior convection coefficients vary with surface orientation and temperature differences, while exterior convection accounts for wind effects, introducing nonlinearities coupling thermal and airflow dynamics.

Radiative heat exchange. Solar radiation transmits through glazing with angle-dependent coefficients for direct and diffuse components. A portion of incoming radiation is absorbed by the glass itself, while the transmitted fraction enters the zone where it is partially absorbed by interior surfaces and partially reflected. These reflections distribute across surfaces according to view factors and absorptivities, with secondary reflections further redistributing energy throughout the space. Long-wave radiation between internal surfaces depends on view factors and emissivities, creating exchanges that vary with surface temperatures to the fourth power. Solar gains introduce time-varying disturbances from sun position, weather conditions, and building orientation.

Thermal bridges. Thermal bridges represent localized high-conductivity paths that occur at structural discontinuities, typically at corners, junctions between walls and floors, and penetrations through the building envelope. These heat leakage paths bypass the main insulation layers, creating direct thermal connections between interior zone air and exterior ambient conditions. In RC network terms, thermal bridges appear as additional resistance branches connecting the zone thermal capacitance directly to outdoor temperature, parallel to the main envelope resistance path. While often neglected in simplified models, thermal bridges significantly affect energy performance, particularly in well-insulated buildings where they can account for a substantial fraction of total heat losses.

Inter-zone air exchange. Inter-zone air exchange uses pressure-driven flow equations coupling zone temperatures, pressures, and door states. Bi-directional airflow depends nonlinearly on temperature and pressure differences, with flow rates determined by geometry and pressure coefficients, creating thermal interactions that evolve with occupancy and door usage. The combination of these multi-physics, multi-scale, and nonlinear phenomena results in a high-fidelity building model that presents realistic challenges for control algorithm validation.

The models include detailed representations of HVAC equipment dynamics that introduce multiple time scales and nonlinearities. Fluid dynamics in hydronic systems capture flow-pressure relationships through pump curves, hydraulic resistances, and pressure-dependent flow distribution across zones. Heat exchanger models employ nonlinear dynamics such as effectiveness-Number of Transfer Units

or Log Mean Temperature Difference methods to represent the coupling between fluid flow rates, temperatures, and heat transfer rates. Air-based systems model mixing processes in economizers, air distribution networks, and fan performance curves that exhibit nonlinear power-flow relationships.

Equipment thermal inertia introduces asymmetric power response characteristics that complicate control design. Hydronic radiators exemplify this challenge: power output increases rapidly by admitting hot water through valve control governed by fluid dynamics, but power reduction requires convective cooling of the thermal mass, a substantially slower process. This asymmetry creates control difficulties at both hierarchical levels. At the tracking level, prolonged saturation during cooling phases leads to integral windup in PI or PID controllers, degrading power reference tracking performance. At the supervisory level, when MPC commands abrupt power reductions to exploit load flexibility, the physical system cannot respond immediately due to thermal inertia. The resulting delay between commanded and actual power delivery causes energy waste and undermines the economic optimization objective. These transient responses span time scales from seconds for valve actuation to hours for complete thermal mass cooling, requiring careful consideration in both model formulation and control design.

Control actuators exhibit realistic nonlinear behaviors including valve flow characteristics, pump speed-flow relationships, and damper position-flow curves. Actuator dynamics incorporate time delays, rate limits, and saturation effects. Sensor models include measurement delays, quantization, and noise characteristics that affect controller performance. These realistic actuator and sensor representations ensure that control strategies validated in BOPTEST face challenges similar to real implementations.

While BOPTEST captures essential building thermal complexities, certain idealizations are introduced to balance model fidelity with computational efficiency. First, sensor measurements are noise-free, eliminating the need for signal filtering and state observation techniques. This idealization isolates control algorithm performance from measurement noise handling, enabling direct assessment of thermal management strategies without confounding effects from sensor uncertainties. Second, building envelope thermal dynamics employ multi-layer RC networks that discretize walls, roofs, and floors into manageable state-space representations. Each envelope component uses multiple thermal resistance and capacitance nodes to capture spatial temperature gradients and transient heat storage, while other physical processes such as radiative exchange, convection correlations, and solar heat gains retain their nonlinear formulations. This approach preserves dominant thermal mass effects and multi-physics coupling essential for realistic control evaluation, while maintaining computational tractability for extensive controller testing across diverse operating scenarios. Third, fluid dynamics in hydronic and air distribution systems are quasi-static. Pipeline pressure dynamics, valve response transients, and flow propagation delays are omitted, with the assumption that fluid transport reaches steady-state instantaneously relative to thermal time scales. This simplification decouples fast hydraulic transients from slower thermal dynamics, reducing model computational burden. These idealizations maintain the nonlinear, multi-zone, multi-physics character essential for control validation while ensuring simulation feasibility for systematic controller comparison and performance benchmarking.

Indoor air humidity plays a crucial role in thermal comfort experience, particularly during cooling operations. Humidity levels directly affect occupant perception of thermal conditions through evaporative cooling from skin, with high humidity reducing the body's ability to dissipate heat and amplifying discomfort at elevated temperatures. The psychrometric properties of air (relating temperature, humidity, and enthalpy) are fundamental to cooling system design and control, where dehumidification often constitutes a significant portion of cooling loads. However, the BOPTEST test cases employed in this study do not provide measurements of indoor air humidity. Humidity control

requires dedicated sensors, actuators, and control strategies that differ substantially from heating management. This research therefore focuses exclusively on heating scenarios.

2.3 Multizone Test Cases: MZRH and MOSA

This thesis use two test cases from BOPTEST that represent distinct building types and HVAC configurations. The test case MZRH and MOSA span the diversity of HVAC equipment types commonly encountered in building thermal management, confirming the generalizability of the proposed hierarchical control framework across different equipment configurations and operational requirements.

Multi-zone Residential Hydraulic (MZRH) Test case MZRH is used for Chatper 2 case study. The

MZRH test case represents a residential building designed according to French thermal regulation RT2012, which establishes requirements for building energy performance. The building has eight zones: living room, three bedrooms, bathroom, hallway, garage, and an attic above all zones, as shown in figure I.4. The total area of the residential building is 81.08 m^2. Five zones are equipped with hydraulic radiators controlled by thermostatic valves. The hallway radiator operates without local control to balance distribution circuit pressure.

The hydronic heating system includes a gas boiler providing hot water, a circulation pump controlling total flow rate, and a three-way mixing valve that regulates supply temperature by blending boiler output with return water. Each radiator exhibits significant thermal inertia, with heating power increase achieved through active heat transfer from hot water circulation, while power decrease relies on passive convective cooling to the indoor environment. This asymmetry in power response—where heating capabilities significantly exceed cooling rates—presents a particular challenge for predictive control that must respect these physical constraints.

The building simulation uses historical weather data from Bordeaux, France (2014), providing realistic outdoor temperature variations (-5°C to 20°C) and solar radiation patterns. The test case includes occupancy schedules with associated internal heat gains, natural ventilation through infiltration, and thermal coupling between adjacent zones. Control objectives include minimizing energy consumption while maintaining zone temperatures above 19°C during occupied periods, balancing comfort requirements with operational costs.

The MZRH test case provides a baseline control solution including two parts. For thermal management, decentralized PI controllers regulate radiator valves in each zone to track temperature references, which is set at the lower comfort boundary. For the supply system, a standalone control design secures the safety of the primary energy source by preventing boiler overload. This supply system controls manipulates the three-way mixing valve that blends boiler output with return water. It also control the circulation pump that maintains the flow rate. These coordinated actions maintain stable and quasi-constant hot water supply temperature for the radiator network.

In this work, validation on the MZRH test case retains the supply system control solution from the baseline strategy. The proposed hierarchical control schemes replace the zone-level thermal management layer. The evaluation focuses on multi-zone coordination, thermal comfort maintenance, and energy efficiency improvements by predictive control at the zone level.

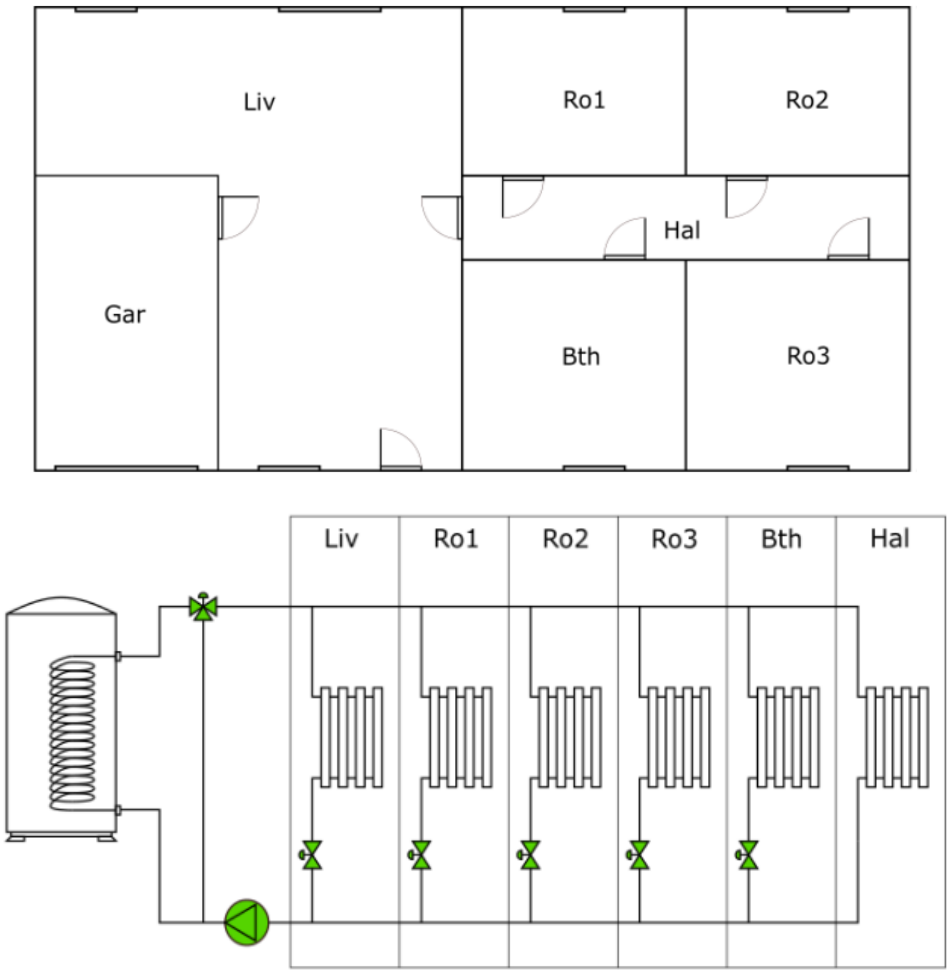

Figure I.3: Adjacency of the MZRH test case and hydronic circuit of heating system.

The control input is

$$\mathbf{v} = [v_1, ..., v_5], \tag{3}$$

where v is the valve opening percentage of radiator. $\mathbf{v}$ is the vector of five zones' valve input.

The measurement vectors are define as

$$\mathbf{P} = [P_1, ..., P_6], \tag{4}$$
$$\mathbf{T} = [T_1, ..., T_8], \tag{5}$$
$$\mathbf{d} = [G, T_{ext}], \tag{6}$$

where $\mathbf{P}$ is a vector of six zones' thermal power input in $[W]$. $\mathbf{T}$ is a vector of all eight zones air temperature in $[K]$. $\mathbf{d}$ is a vector of weather condition, including solar ratiation G in $[W/m^2]$ and exterior air temperature T_{ext} in $[K]$.

Multi-zone Office Simple Air (MOSA) MOSA validates the methodology in chapter III and IV.

The MOSA test case simulates the middle floor of a medium office building, which is derived from U.S. Department of Energy Prototype Building Models [48]. The floor is divided into five zones:

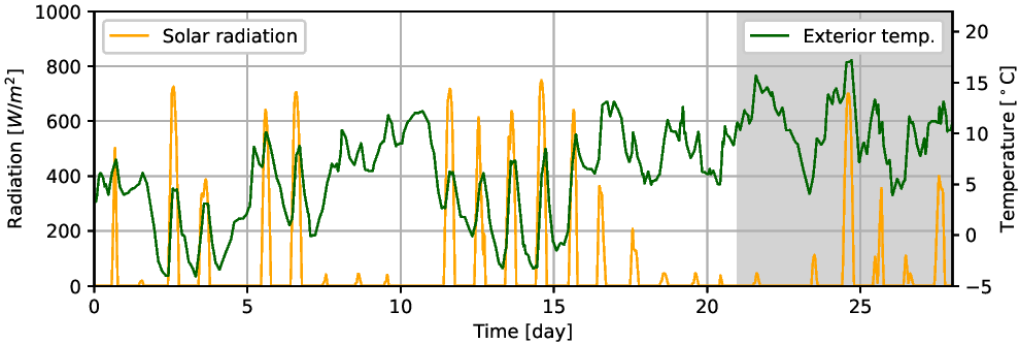

Figure I.4: MZRH test case radiation and exterior temperature profile.

one interior (central) zone and four perimeter zones (north, east, south, west) exposed to external environmental conditions, as shown in figure I.5. This configuration represents typical commercial building layouts where perimeter zones experience significant solar gains and envelope heat losses, while interior zones are dominated by internal gains.

The HVAC system employs an Air Handling Unit (AHU) with Variable Air Volume (VAV) terminal units for zone-level control, following American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) standard sequences of operation. Outside fresh air and return air are mixed in economizers before entering the AHU heating coil for primary heating. Conditioned air is distributed through dampers to individual VAV units, each equipped with reheat coils for secondary zone-level temperature control.

The system operates in two modes based on occupancy schedule. Control objectives focus on maintaining zone temperatures above 19 $°C$ during occupied periods (9:00-17:00) while minimizing total energy consumption across both AHU and VAV systems. The ventilation system provide controlled airflow $q_{air,VAV}$ (almost a constant flow) to VAV and VAV reheat functions are enabled for individual zone control. During unoccupied periods (17:00-9:00), the ventilation system operates with reduced airflow and heating is provided solely through the AHU, with VAV reheat functions disabled.

The supply system control manages the AHU and ventilation to ensure reliable operation. The economizer dampers mix outdoor fresh air with return air before entering the AHU heating coil for primary heating. This centralized heating provides the baseline thermal input to all zones.

The MOSA test case incorporates multiple nonlinear elements including valve characteristics for water-to-air heat transfer in heating coils, mixing damper dynamics, and variable-speed fan operation. Thermal management faces disturbances from occupancy variations, solar radiation with time-dependent incident angles on different façades, and outdoor weather conditions. The MOSA case validates control performance on air-based HVAC systems with centralized-decentralized architecture, complementing the MZRH validation on hydraulic systems.

In this work, validation on the MOSA test case retains the supply system control solution from the baseline strategy, including AHU heating and ventilation management with fixed airflow settings. The proposed hierarchical control schemes replace only the zone-level VAV reheat control during occupied periods from 9:00 to 17:00. The evaluation focuses on multi-zone coordination, thermal comfort maintenance during occupied hours, and energy efficiency improvements by predictive control at the zone level.

The control input is

$$\mathbf{v} = [v_1, ..., v_5], \tag{7}$$

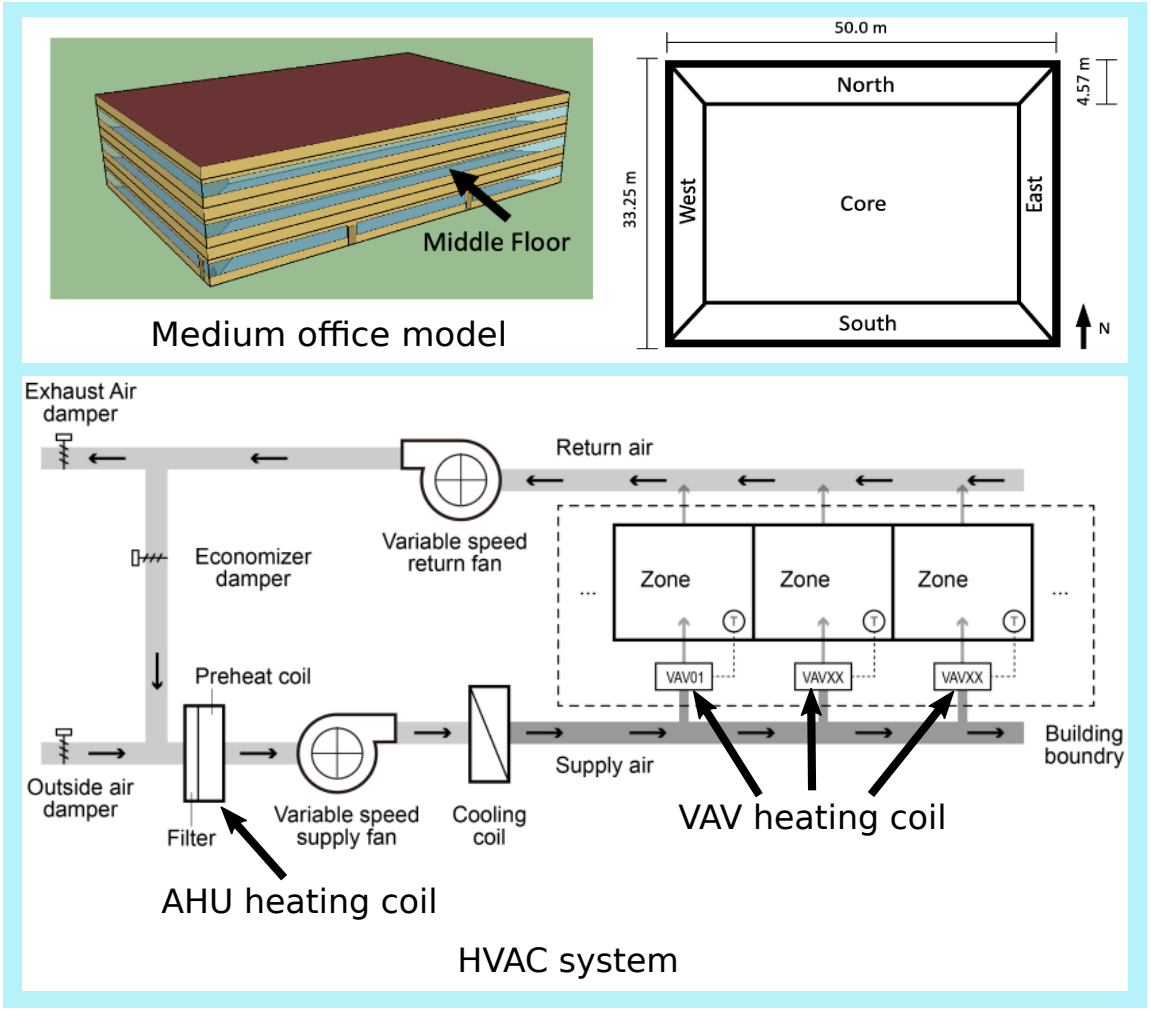

Figure I.5: BOPTEST benchmark Multi-zone Office Simple Air test case.

where v is the valve opening percentage of heating coil. $\mathbf{v}$ is the vector of five zones' valve input.

The measurements in test case are thermal power of VAV $\mathbf{P}$, zone temperature $\mathbf{T}$ and weather condition $\mathbf{d}$. Vectors are define as

$$\mathbf{P} = [P_1, ..., P_5], \tag{8a}$$
$$\mathbf{T} = [T_1, ..., T_5], \tag{8b}$$
$$\mathbf{d} = [G, T_{ext}], \tag{8c}$$

where the notation is similar to MZRH test case.

2.4 Multi-objective Formulation and Performance Metrics

The control problem for building thermal management is formulated as a multi-objective optimization that balances energy consumption with occupant thermal comfort. The optimal trade-off between these objectives depends on building usage, occupant preferences, and economic considerations.

Energy consumption objective: The energy consumption objective is defined as the total delivered thermal power integrated over the evaluation period:

$$J_{\text{energy}} = \int_{t_0}^{t_f} \sum_{i=1}^{n_z} P_i(t)\, dt \tag{9}$$

where $P_i(t)$ represents the thermal power delivered to zone i, n_z is the number of zones, and $[t_0, t_f]$ defines the evaluation period. This formulation focuses on delivered thermal energy rather than primary energy consumption. Equipment efficiency factors and energy pricing dynamics are excluded. This choice enables fair comparison across different HVAC system types and isolates thermal management performance from equipment-specific characteristics.

For practical applications, this thermal power metric can be converted to fuel consumption (for gas boilers) or electrical consumption (for heat pumps and electric heating) by incorporating equipment efficiency curves and operational characteristics. Furthermore, time-varying energy prices could be integrated into the objective function through a weighting factor $c(t)$. It enables MPC to shift heating loads to periods with lower energy costs while respecting comfort constraints.

Thermal comfort objective: Thermal comfort is quantified through violations of temperature boundaries that define acceptable thermal conditions for occupants. The discomfort metric is formulated as:

$$J_{\text{discomfort}} = \int_{t_0}^{t_f} \sum_{i=1}^{n_z} \max(0, \underline{T}(t) - T_i(t))\, dt \tag{10}$$

where $T_i(t)$ is the measured temperature in zone i, and $\underline{T}(t)$ is the time-varying lower comfort boundary. Upper temperature bounds are not enforced in heating scenarios, as temperatures above the comfort range do not constitute discomfort during winter operations. The integration in the two criteria (9) and (10) is approximated by a sum over the simulation time steps and divided by sampling time.

The overall control objective combines energy consumption and thermal comfort through a weighted sum:

$$J_{\text{total}} = J_{\text{energy}} + \lambda J_{\text{discomfort}} \tag{11}$$

where $\lambda > 0$ determines the relative importance of comfort versus energy savings. Different λ values gives different operating points along the Pareto frontier: low values prioritize energy savings, while high values prioritize comfort. Rather than prescribing a single weighting, results are presented with a range of λ values, illustrating the Pareto frontier for each control approach.

3 Hierarchical Control Methodology

This thesis proposes a hierarchical control framework for building thermal management that addresses the equipment diversity and data limitation challenges identified in the previous sections. The methodology separates concerns between long-term thermal planning and short-term equipment

regulation, enabling generalized deployment across diverse HVAC configurations while maintaining robustness under limited data conditions.

3.1 Hierarchical Control Architecture

The proposed hierarchical control framework constructs a nested closed-loop system that manages building thermal systems through power signals. This separation has been successfully employed in in equipment-specific solutions [49, 12, 33, 35]. This work aims to extend the system segmentation advantages of this scheme to a more scalable control design framework. The two levels correspond to nested control loops in the overall control block diagram, as illustrated in Figure I.6.

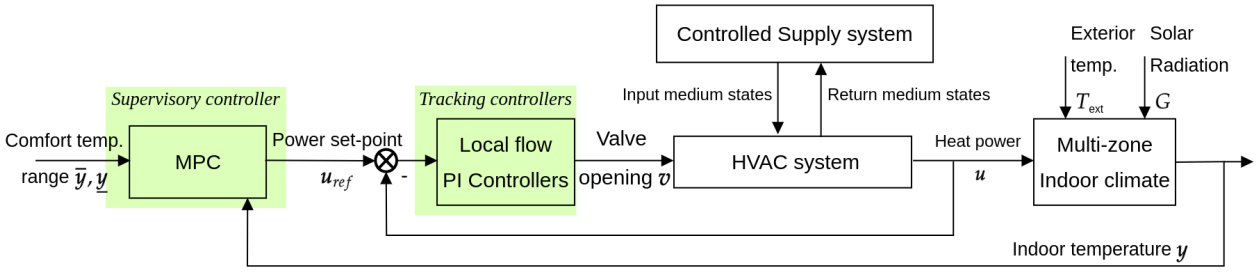

Figure I.6: Hierarchical control scheme with supervisory-level predictive controller and tracking-level controllers. The supervisory controller performs prediction over horizon H with time step T_{sup}, updating thermal power references at each step. During each supervisory interval, discrete PI controllers regulate actuators to track power references with time step T_{track}. Medium refer to thermal mass (like air and water) that circulate enthalpy in equipment.

In the outer loop, the supervisory level uses predictive control to handle economic objectives and indoor climate dynamics. The supervisory controller predicts temperature evolution over the prediction horizon and generates thermal power reference signals u_{ref} for the tracking level. This long-term optimization addresses slow thermal dynamics and system-wide objectives, operating with time step T_{sup} typically in the range of 10-30 minutes. The supervisory controller optimizes energy distribution across building zones to balance thermal comfort and energy consumption.

In the inner loop, the tracking level employs PI controllers to respond rapidly to power demands and regulate actuator components such as valve openings, pump flows, damper positions, or turbines. Each zone employs a dedicated controller that manages equipment-specific variables, operating at a faster time step T_{track} typically in the range of 0.5-5 minutes. These controllers address fast dynamics and equipment constraints while maintaining the power output demanded by the supervisory controller during each supervisory interval.

The hierarchical separation through power signals provides several advantages. It decouples the thermal planning problem from equipment-specific dynamics, allowing the supervisory controller to focus on building-level optimization without detailed HVAC models. It enables equipment-agnostic design at the supervisory level, where the same control formulation applies across different HVAC configurations. It facilitates independent design and tuning of the two control levels, as demonstrated in the implementation workflow. This architecture splits the building system into two distinct components: HVAC system and multi-zone indoor climate. The tracking level handles HVAC system nonlinearities

and disturbances, guaranteeing that the process between power reference given by the supervisory controller and thermal power input to indoor climate remains tractable for predictive control.

3.2 Implementation Workflow and Design Procedure

The implementation of the hierarchical control scheme follows a workflow progressing from data acquisition through controller design to validation, as shown in Figure I.7. The workflow separates tracking level and supervisory level development, with each level having distinct data requirements, design procedures, and validation steps.

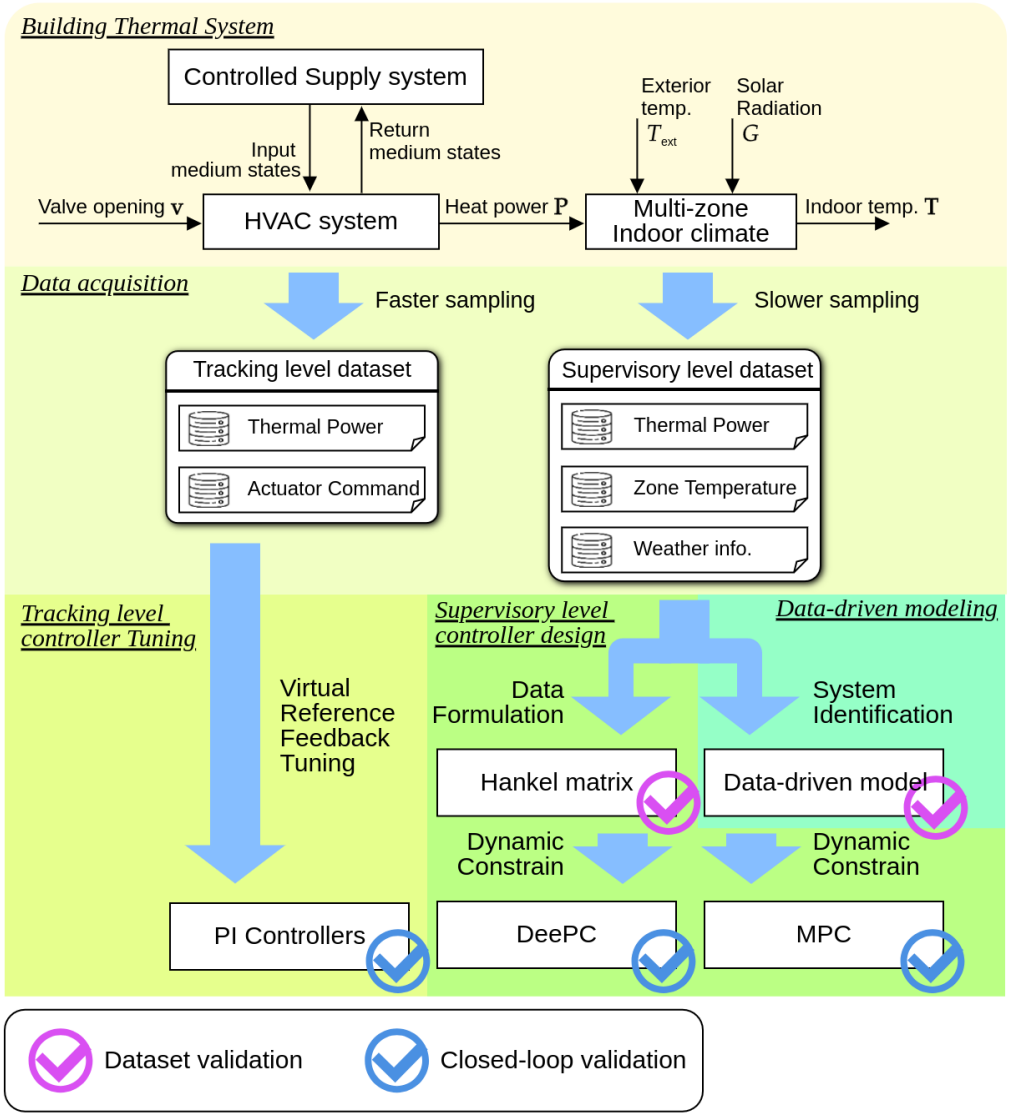

Figure I.7: Implementation workflow of the proposed methodology. Data acquisition collects datasets at different sampling rates for tracking and supervisory levels. Controller design proceeds independently for both levels. Validation includes dataset validation for prediction models and closed-loop validation for complete controllers.

Data Acquisition

Data acquisition collects measurements from the building thermal system at two different time scales. The tracking level dataset employs faster sampling to capture HVAC system dynamics. It includes thermal power measurements and actuator actions. These measurements are obtained through dedicated excitation experiments. Actuators are manipulated using pseudo-random binary sequence (PRBS) signals according to [30]. PRBS excitation provides sufficient frequency content to characterize HVAC power response across the operating range. This experimental approach is practical because HVAC power response experiments are relatively straightforward and can be conducted without compromising occupant comfort.

The supervisory level dataset uses slower sampling aligned with thermal dynamics of the building. It includes thermal power delivered to each zone, zone temperature measurements, and weather disturbances including exterior temperature and solar radiation. This data can be collected from normal building operation or through dedicated thermal experiments. The data acquisition in this level focuses on capturing temperature response at frequencies well below the HVAC power response bandwidth. When designing thermal experiments for informative datasets, actuator inputs are manipulated as square-wave signals that achieve quasi-steady-state power levels while exciting unsteady temperature behavior. The key objective is to excitet temperature variations by inputting randomized power. Accurate power tracking is then not necessary. The development of the two control levels are independent, allowing parallel design workflows.

This data acquisition strategy is based on dynamic mode decomposition principles, which match the sampling rate to the time scales of the dominant dynamic modes of interest [50, 51]. This excitation strategy targets the frequency spectrum between the characteristic frequencies of temperature dynamics and power dynamics, ensuring that thermal response is captured while power transients have settled. Slow thermal modes is then separated from fast power dynamics. Also, the manipulated inputs should vary amplitude and duration to explore different operating points and initial temperature conditions.

Tracking Level Controller Tuning

Tracking level controller design uses VRFT to optimize PI controller parameters. VRFT formulates a data-driven optimization problem that tunes controller parameters to match desired closed-loop reference dynamics without requiring explicit system model. Using the collected actuator action and thermal power data, VRFT determines proportional, integral and anti-windup gains that minimize tracking error with respect to the reference model. The desired reference model is specified with time constants reflecting the required response speed for tracking supervisory-level power commands.

The tuned controllers is validated in closed-loop to assess power tracking performance. Classic control metrics are evaluated including response time, settling time, overshoot, and steady-state tracking accuracy. Step response tests verify that controllers can follow power reference changes within required time frames. The validation confirms that tracking dynamics are sufficiently fast and accurate to support supervisory level control. This ensures that power references from the upper level can be reliably realized by HVAC actuators.

Supervisory Level Controller Design

Supervisory level controller design begins with data-driven modeling or formulation using the collected temperature and power datasets. Two parallel approaches are investigated. For model-based predictive control, system identification produces data-driven models representing building thermal

dynamics. Gray-box modeling using equivalent RC networks translates building thermodynamics through physical insight and parameter estimation. Black-box modeling using 4SID derives state-space representations directly from input-output data without physical assumptions. For model-free approaches, DeePC DeePCconstructs Hankel matrix from historical input-output trajectories, encoding system behavior implicitly without explicit model extraction.

Dataset validation assesses model or formulation quality through prediction testing. The training dataset is selected from weeks with similar weather conditions to ensure consistency and fairness in model comparison. This selection strategy minimizes the influence of indoor climate's external input on validation results. Prediction accuracy is quantified through mean squared error over prediction horizons. The validation confirms whether identified models or formulated data matrix can accurately predict temperature evolution instead of overfitting the training dataset.

Controller implementation integrates the validated model or Hankel matrix into predictive control as a constrain. Rate constraints on power changes may be incorporated to account for asymmetric HVAC dynamics, with constraint parameters identified from operational data. Multi-objective performance is evaluated through Pareto frontier analysis across different weighting factors. Sensitivity analysis examines controller robustness to hyperparameter variations, including prediction horizon, regularization parameters, and constraint settings. Computational efficiency is evaluated through optimization solution times to ensure real-time feasibility.

Chapter II

Hierarchical MPC with Gray-box and Black-box Models

1 Literature Review on MPC for Buildings 48
1.1 MPC Classification and Implementation Approaches 48
1.2 Gray-box model vs. Black box model for buildings 50
2 System Identification for Prediction Models 54
2.1 Gray-box Resistance-Capacitance Model Identification 54
2.2 Practical Implementation Considerations 61
3 Hierarchical MPC Design 62
3.1 Supervisory-level MPC Formulation 63
3.2 Tracking-level Controller Tuning via VRFT 66
4 Simulation Validation and Performance Evaluation 69
4.1 Open-loop Validation of Indoor Climate Models 69
4.2 Tracking Controller Performance Assessment 71
4.3 Closed-loop Hierarchical Control Performance 72

1 Literature Review on MPC for Buildings

Designing MPC strategies for building temperature control involves three main considerations: the control objectives (i.e., the cost function and constraints), the implementation architecture (i.e., the level of centralization and control hierarchy), and the type of prediction model used within the MPC formulation. This section reviews existing MPC approaches applied to buildings in order to identify the structure of the proposed hierarchical control framework and to situate it within the broader range of MPC implementation strategies.

1.1 MPC Classification and Implementation Approaches

According to control objectives, the cost function design can be divided into two main types: tracking MPC and economic MPC. Tracking MPC aims to keep the system close to a desired reference trajectory while satisfying all system constraints. This approach benefits from well-established stability and optimality guarantees [52, 53]. The objective function penalizes deviations from the reference states, ensuring convergence properties that support both theoretical analysis and practical tuning.

In contrast, economic MPC focuses on optimizing an economic or performance-related objective that may include operational costs, energy efficiency, or other performance metrics beyond simple state regulation [54]. The term “economic” here refers not only to monetary costs but to any performance measure beyond stabilization objectives. Common examples in energy systems include minimizing

energy consumption, participating in demand response programs, or reducing peak loads [6, 7]. Unlike tracking MPC, economic MPC does not inherently guarantee closed-loop stability to a fixed equilibrium [55], since it searches for operating points that minimize the chosen cost function [56]. This approach enables direct integration of operational objectives into the control layer rather than separating economic optimization and dynamic control into deeper hierarchical layers. For building thermal management, this integration allows the controller to anticipate time-varying energy prices and exploit building thermal inertia to reduce costs while maintaining comfort.

From a system architecture perspective, MPC implementations can be classified based on their role within the control hierarchy and their interface with the physical system. This classification reflects how control authority is distributed through system decomposition. Three primary categories are typically identified in the literature: direct actuator controller, tracking reference governor, and Intermediate variable programmer, as illustrated in Figure II.1.

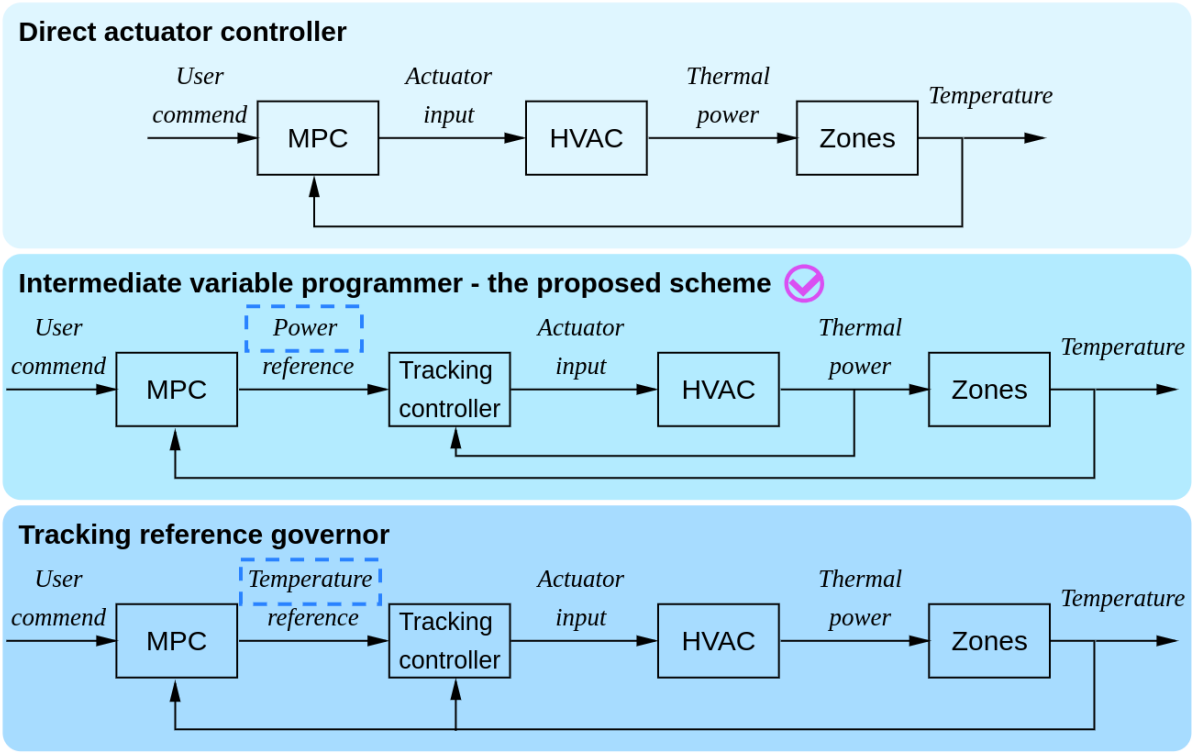

Figure II.1: Classification of MPC implementation paradigms by control interface.

A direct actuator controller issues low-level control commands, such as valve openings, damper positions, or pump speeds, directly within the MPC optimization. This approach requires a detailed system model that captures both the thermal dynamics and actuator behavior. However, challenges arise when system dynamics cover a wide range of time scales—from fast actuator responses to slow building thermal changes. As a result, model accuracy and data quality requirements become very strict [49].

Tracking reference governor formulations assume that lower-level controllers already exist to track state setpoints. For building thermal control, local controllers perform temperature tracking in each zone, while the MPC layer schedules temperature setpoints to optimize economic objectives [12, 57, 58]. This architecture offers advantages in system upgrading and integration, as temperature

tracking services are typically provided by existing equipment controllers in building management systems. The MPC can be added as a supervisory layer without modifying the underlying control infrastructure. From a system segmentation perspective, temperature setpoints create coupling between building thermal dynamics and HVAC equipment dynamics. The supervisory MPC must account for both slow thermal evolution and fast equipment responses when optimizing temperature trajectories, as these two subsystems interact directly through the temperature variable.

Intermediate variable programmer generate reference signals for intermediate physical variables—such as thermal power or energy flow—that bridge high-level optimization goals and low-level actuation. This hierarchical structure closes an inner tracking loop through feedback of intermediate physical quantities, typically representing energy transfer or material flow rates. Lower-level tracking controllers regulate these intermediate variables, effectively isolating part of the system dynamics from the supervisory optimization and reducing its computational load [33]. For building thermal control, the interaction between HVAC equipment and indoor climate occurs through heat power transfer. System segmentation at the power signal level isolates indoor climate dynamics from equipment-specific complexity, addressing equipment diversity by abstracting away actuator details such as valve characteristics, damper positions, or compressor dynamics. However, this approach requires measurement or estimation of intermediate variables like thermal power, which are not always directly available from standard building sensors and may necessitate additional instrumentation or soft-sensing techniques.

Model-free approaches offer a path that identifies tracking controller parameters directly from data without requiring explicit system models. Virtual Reference Feedback Tuning (VRFT) [40] is one such data-driven method that has been deployed in various engineering applications [59, 41, 60, 61]. Recent work on water reservoir control [62] demonstrated this methodology in managing complex, nonlinear systems. This thesis employs VRFT within the proposed control hierarchy to address equipment diversity in building thermal management. Using input-output data, VRFT tunes controller parameters to match a desired closed-loop reference model. This approach accommodates diverse HVAC systems at low modeling cost and is particularly suitable for simple, analyzable controller classes such as PID. This work adapts VRFT to building thermal management, enabling equipment-independent control design across different HVAC configurations.

The methodology proposed in this study follows the Intermediate variable programmer paradigm, as highlighted in Figure II.1. The supervisory MPC optimizes thermal power references, while the tracking-level controllers handle the specific dynamics of building equipment. The detailed implementation of this hierarchical architecture is presented in Figure I.6 in Chapter 1.

1.2 Gray-box model vs. Black box model for buildings

The architectural choices discussed above determine how MPC interfaces with the physical system. Complementing this, the prediction model determines how the controller anticipates future system behavior within the optimization. The choice between gray-box and black-box modeling approaches shapes the data requirements, generalization capabilities, and reliability of predictive models in control applications. Gray-box models incorporate physical structure and domain knowledge into the model framework, constraining the parameter space through conservation laws and known relationships. Black-box models learn system dynamics entirely from measured data without explicit physical assumptions, relying on flexible mathematical structures to approximate input-output mappings.

This section examines the trade-offs between these paradigms, with emphasis on their robustness in extrapolation and susceptibility to overfitting—critical considerations when deploying data-driven control under limited or imperfect data conditions.

Gray-box modeling

Gray-box modeling can be realized through two approaches. The first approach uses constraints that preserve physical information during model parameterization. For instance, in models with graph structure, the adjacency matrix can be incorporated as a physical constraint within the parameter identification optimization, as shown in RC network identification [35] and Neural ODE formulations [63]. This constraint enforces sparsity in the coupling relationships of internal dynamics, reducing the degrees of freedom. The second approach augments the loss function with regularization terms that balance fitting physical equations against fitting data, as exemplified by Physics-Informed Neural Networks (PINNs) [64]. This formulation prevents overfitting by penalizing solutions that violate known physical laws.

Physical constraint-based regularization operates differently from general mathematical penalization: physical constraints eliminate parameter space regions that violate fundamental laws [65]. In thermal system modeling, energy conservation requirements bound the parameter space, preventing models from predicting physically impossible behaviors such as spontaneous temperature increases without heat input or violations of thermodynamic equilibrium. The structural constraints in gray-box formulations act as implicit priors that guide parameter estimation toward physically meaningful solutions, particularly when training data are limited or incompletely cover the operational envelope [66, 67]. Conversely, black-box models achieve flexibility through high-dimensional parameter spaces that can approximate arbitrary input-output mappings, but this requires extensive data coverage [30, 68].

The application of gray-box modeling for building thermal control has been extensively studied across various contexts. Studies in chemical process control [69], environmental modeling [70], and fluid dynamics [64] demonstrate that physics-informed models maintain prediction reliability during extrapolation through structural constraints. For building applications specifically, RC network-based gray-box models have been applied to demand response control [71], retrofit design analysis [72], and multi-zone thermal management [73].

Black-box modeling

Black-box modeling approaches learn system dynamics purely from measured data without imposing physical structure. Representative black-box methods include subspace identification techniques such as Numerical algorithms for Subspace State Space System IDentification (N4SID) and Multivariable Output-Error State sPace (MOESP) algorithms, which extract state-space representations directly from input-output measurements through projection-based methods [16, 74]. Direct data-driven frameworks based on Willems' fundamental lemma bypass explicit model identification by utilizing Hankel matrices constructed from historical trajectories for prediction [45]. Regression approaches, including autoregressive models with exogenous inputs (ARX) and nonlinear autoregressive models (NARX), capture dynamics through polynomial or nonlinear function fitting [30]. Machine learning methods, ranging from neural networks and support vector machines to Gaussian processes and ensemble methods, provide flexible representations capable of approximating arbitrary input-output mappings [75, 76].

Machine learning techniques have enabled purely data-driven modeling approaches for building thermal systems. Neural network architectures, including Long Short-Term Memory networks and deep feedforward networks, have demonstrated strong performance for building load prediction and temperature forecasting when trained on large datasets [77, 78]. However, their application in MPC faces challenges related to non-convex optimization landscapes and limited interpretability, motivating recent developments in convex neural network architectures such as Input Convex Neural Networks [79, 80]. This thesis focuses on reducing modeling and tuning burden for practical deployment in existing buildings with limited historical data. Therefore, machine learning approaches requiring extensive datasets and specialized expertise are not deeply investigated here. Instead, the study emphasizes linear state-space models obtained through subspace identification, which offer computational tractability, theoretical maturity, and straightforward integration into standard MPC formulations.

Discussion on the bias-variance tradeoff

The theoretical foundation for comparing gray-box and black-box modeling performance rests on the bias-variance decomposition of prediction error [81, 82]. For a prediction model $\hat{f}(x)$ trained to approximate the true function $f(x)$, the expected prediction error on unseen data decomposes as:

$$\mathbb{E}[(y - \hat{f}(x))^2] = \text{Bias}^2[\hat{f}(x)] + \text{Var}[\hat{f}(x)] + \sigma^2 \tag{12}$$

where

$$\text{Bias}[\hat{f}(x)] = \mathbb{E}[\hat{f}(x)] - f(x), \quad \text{Var}[\hat{f}(x)] = \mathbb{E}\left[(\hat{f}(x) - \mathbb{E}[\hat{f}(x)])^2\right] \tag{13}$$

and σ^2 represents irreducible measurement noise [75]. In dynamic system identification, x represents past inputs u and outputs y used to predict the next system output. The objective is to find $\hat{f}$ that captures the system dynamics and generalizes well to new operating conditions. This requires managing the bias-variance trade-off:

- **Bias** quantifies systematic errors from structural mismatches between the model and the true system. In building thermal modeling, bias arises when model assumptions neglect physical phenomena—for instance, a fully connected thermal network that ignores the actual zone adjacency structure, or a linearized formulation that does not capture nonlinear heat transfer. These errors persist regardless of training data size, leading to underfitting.
- **Variance** quantifies the model's sensitivity to the specific training dataset used for parameter estimation. It arises from fitting model parameters to noisy measurements, causing the estimated model to differ across different data samples from the same system. High variance indicates excessive adaptation to sample-specific noise rather than generalizable patterns, also known as overfitting.

Gray-box and black-box models occupy different positions in the bias-variance trade-off. Gray-box models impose physical structure that restricts the hypothesis space, introducing inductive bias through conservation laws and known relationships. This structural constraint may increase bias when physical assumptions mismatch the true system, but reduces variance by preventing overfitting to training data noise [66]. Recent studies on building thermal systems demonstrate that gray-box RC models can achieve considerable variance reductions at the cost of bigger bias, particularly when dealing with unmeasured disturbances such as occupant loads [83].

Black-box models employ flexible hypothesis spaces with minimal structural assumptions, allowing them to capture complex nonlinear dynamics and subtle coupling effects that simplified physical models might miss [30]. This flexibility reduces bias but increases variance, as the model becomes sensitive to training data characteristics and requires larger datasets to constrain parameter estimates effectively [81]. Without sufficient data coverage, black-box models risk overfitting and produce unreliable predictions when extrapolating beyond the training distribution. Recent comparative studies on building thermal systems confirm this trade-off: neural network models demonstrate superior average prediction performance compared to RC gray-box models when abundant training data are available, though both approaches are affected by building characteristics, seasonal variations, and data quality [78, 84]. The practical implication for building thermal control is that gray-box models offer superior generalization under limited data through physics-informed regularization, while black-box approaches excel when abundant diverse data cover the operational envelope comprehensively [85, 86].

For MPC applications, modeling objectives extend beyond prediction accuracy to ensure optimality of control inputs within the prediction horizon [87, 88]. The optimization relies on model gradients to guide the search for optimal trajectories. However, models are calibrated only at operating points present in identification datasets, requiring extrapolation to explore state-space regions where optimal solutions may reside. Under closed-loop data collection, system states naturally concentrate around limited operating points dictated by feedback control [89], creating a challenge when MPC optimization identifies solutions in regions where model extrapolation is unreliable. This extrapolation requirement makes the bias-variance trade-off particularly critical for MPC: gray-box models maintain reliability through physical constraints that remain valid beyond training conditions [69, 90], while black-box models may produce predictions that violate physical principles when input distributions shift beyond training regimes [75, 91].

1.3 Research Questions and Chapter Organization

This chapter addresses the research question: how do gray-box and black-box modeling approaches perform differently in MPC under limited historical operational data? The bias-variance framework suggests that gray-box models trade flexibility for physics-informed regularization, while black-box models maximize representational capacity at the cost of higher sensitivity to training data. Although both approaches may achieve comparable open-loop prediction accuracy, their performance can diverge during closed-loop operation, when the MPC optimization explores state-space regions distant from training conditions. This study investigates whether gray-box models achieve superior closed-loop control performance through reliable extrapolation, even at the cost of higher open-loop prediction error. The hierarchical control framework enables this comparison by separating thermal dynamics modeling at the supervisory level from equipment regulation at the tracking level, where model-free VRFT handles actuator diversity without requiring detailed HVAC system models.

Section 2 presents gray-box modeling using RC networks, translating building thermodynamics into electrical circuit topology with convex parameter identification. Section 3 formulates the economic MPC problem minimizing energy consumption under thermal comfort constraints, with both modeling approaches integrated into the prediction horizon. Subsection 3.1 addresses tracking-level control using VRFT for model-free PID tuning to achieve power reference tracking. Section 4 presents validation using BOPTEST testbeds, comparing open-loop prediction accuracy and closed-loop control performance across training dataset lengths. Results demonstrate that despite comparable

validation performance, gray-box models achieve superior closed-loop energy efficiency and comfort maintenance through reliable extrapolation enabled by physics-informed structural constraints.

2 System Identification for Prediction Models

2.1 Gray-box Resistance-Capacitance Model Identification

Gray-box modeling reduces reliance on large amounts of data by incorporating prior knowledge of the model structure [6, 31, 32]. The RC model is a representative gray-box approach that translates building thermodynamics into electrical circuit topology. Its capability for modeling indoor climate has been extensively studied [33, 34, 35, 36]. This section presents the gray-box identification methodology adopted in this work, beginning with the lumped modeling framework that underlies RC network construction, followed by specific model structures and parameter estimation procedures.

2.1.1 RC Network Analogy for Building Thermodynamics

The transition from building thermodynamics to RC networks requires a lumped modeling approach that reduces spatial complexity while preserving essential dynamic characteristics. Lumped modeling is a simplification technique used in building thermal analysis to reduce the complexity of heat transfer calculations while maintaining sufficient accuracy for control application. In a lumped model, the building or its components are represented as a network of nodes, each characterized by a uniform temperature. These nodes are connected by thermal resistance and capacitance, analogous to electrical circuits. The key assumptions of lumped modeling include:

- Spatial uniformity: Each lumped element (e.g., a wall or air) is assumed to have a uniform temperature throughout its volume. This implies that internal temperature gradients within each element are ignored.
- Concentrated thermal mass: The thermal capacitance of each element is concentrated at a single point (node) rather than distributed throughout its volume.
- Linear heat transfer: Heat transfer between nodes is assumed to be linear and proportional to the temperature difference, allowing the use of thermal resistances.
- Constant thermal properties: Material properties such as thermal conductivity and specific heat capacity are assumed to be constant and independent of temperature.

These simplifications allow for the creation of tractable mathematical models that can be solved analytically or simulated numerically with relatively low computational cost. While lumped models sacrifice some level of detail compared to distributed parameter models, they often provide sufficient accuracy for building-level energy analysis and control strategy development.

The RC analogy provides a powerful conceptual and mathematical tool for implementing lumped models of building thermal systems. This approach draws a parallel between thermal systems and electrical circuits, where:

- Thermal resistance (R) represents the opposition to heat flow, analogous to electrical resistance. It is measured in K/W (Kelvin per Watt) and is typically associated with conduction through materials or convection and radiation at surfaces.
- Thermal capacitance (C) represents the ability to store thermal energy, analogous to electrical capacitance. It is measured in J/K (Joules per Kelvin) and is related to the mass and specific heat capacity of materials.
- Temperature difference (ΔT) is analogous to voltage difference in electrical circuits.
- Heat flow rate (ϕ) is analogous to electrical current.

This RC framework aligns with bond-graph theory, which represents dynamic systems through three fundamental physical phenomena. Thermal resistance capture dissipative behavior (energy loss), thermal capacitance capture storage of potential energy (heat accumulation), and inertial effects capture dynamic momentum. Building thermal systems are dominated by dissipative and storage phenomena, making RC networks well-suited for their representation.

Using this analogy, building elements can be represented as networks of resistors and capacitors, enabling the application of circuit analysis techniques to thermal problems. This framework facilitates intuitive model construction and supports parameter identification from measurement data.

2.1.2 First-Order RC Model Identification

Extraction of the Gray Box Model

The first-order RC model provides a simplified yet effective representation of building thermal dynamics. This section presents the model structure in continuous-time form, establishes the system identification approach, then describes the discretization needed for parameter estimation and control implementation. The single-zone identification formulation is presented first to establish concepts, then extends to multi-zone configurations that capture inter-zone thermal coupling.

The single zone is modeled as a first-order RC circuit where the entire thermal mass and behavior are concentrated into one element, neglecting internal heat transfers and assuming a uniform temperature distribution. In continuous-time, the thermal behavior is described by the following differential equation:

$$C_i\dot{T}_i = \sum_{\forall j \in N_{ad,i}} \frac{1}{R_{i,j}}(T_j - T_i) + \frac{1}{R_{i,ex}}(T_{ext} - T_i) + P_i + S_iG \tag{14}$$

This equation is structurally similar to nodal analysis in electrical circuit theory. $N_{ad,i}$ is a set of neighbor zones. The equation describes thermal exchange between the zone and its surrounding environment through the thermal resistance $R_{i,ex}$, along with internal heat sources P_i and solar radiation gains S_iG. Note that the time index is implicit throughout: temperatures $T_i(t)$, inputs $P_i(t)$, and disturbances vary continuously with time t.

For a multi-zone thermal system, this work adopts the centralized modeling approach used in [34]. The first-order RC Multi-Zone (1RCMZ) model extends the single-zone representation to capture the indoor climate of a multi-zone building. The multi-zone configuration incorporates the following elements:

- Multiple Zones: The system consists of a set of zones represented as N_{all}. The total number of zones is n. The adjacency is defined as matrix Ad in the size of $(n \times n)$. Each scalar at row i and column j is $Ad_{(i,j)}$

$$Ad_{(i,j)} = \begin{cases} 0 & \text{, if zone i and j not connected, or } i = j, \\ 1 & \text{, if zone i and j connected.} \end{cases} \tag{15}$$

  $N_{ad,i} := \{j | Ad_{(i,j)} = 1, \forall j \in N_{all}\}$ defines a set of neighbor zone index.

- Thermal Exchange: Each zone $i \in N_{all}$ has an equivalent thermal capacitance C_i, a temperature T_i, a thermal power source P_i. $R_{i,j}$ is thermal resistance between each pair of adjacency.

- External Walls and Floor: Assume that all zones i either have external walls or are in contact with the external ground. In equivalent circuit, the connection to exterior are simplified into a constant voltage source T_{ext} and a resistance $R_{i,ex}$ in series.

- Solar Radiation Input: All zones are subject to solar radiation. The radiation intensity is consistent across the system, denoted as G. Each zone has an equivalent incident area for solar radiation labeled S_i.

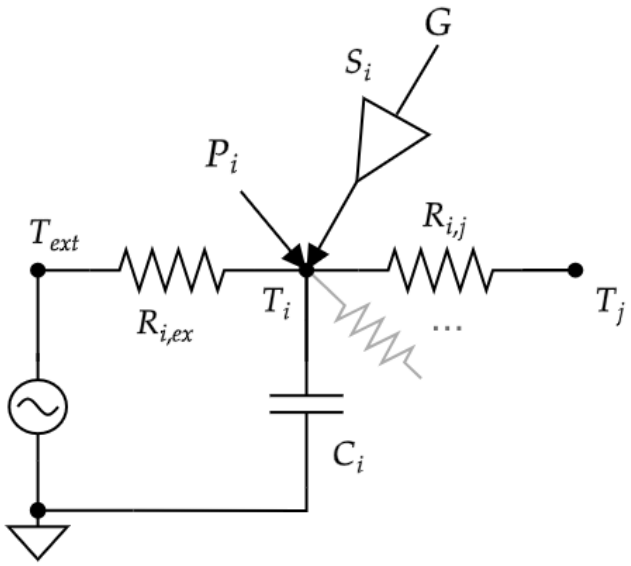

Figure II.2: Equivalent 1st order RC circuit of multi-zone gray-box modeling, showing a sample rooms as an example. Triangle shapes in the diagram represent equivalent surfaces, indicating the gain from solar radiation to absorbed radiative power.

The resulting multi-zone network, illustrated in Figure II.2, allows for the modeling of inter-zone heat transfer and overall building thermal behavior while reducing computational complexity. The system identified by the RC structure has two intrinsic properties:

- Every row in state space model corresponds to a node in thermal graph. Similar to electric circuit, nodal analysis here implies that the sum of heat flow is zero. The multi-zone model respects thus energy conservation when taking interaction into account.

- The constrution adjacency fully defines, describes, the energy distribution in multi-zone.

The distributed expression in equation (14) can be concatenated in a state-space model; figure II.3 summed up finally the RC scheme for a given building.

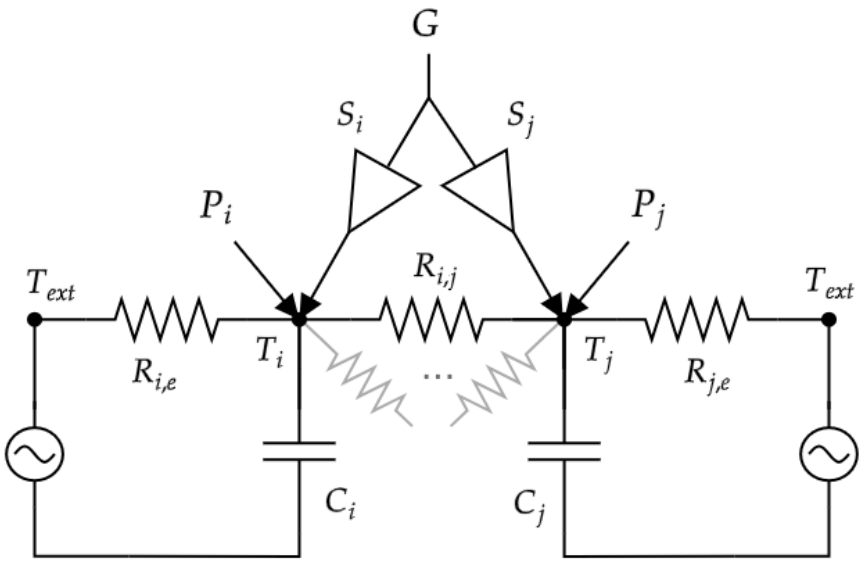

Figure II.3: Multi-zone 1st order RC model.

The formulation below defines input, disturbance, and output vectors, along with the continuous state-space model matrices. The input vector u represents the controllable power sources in each zone:

$$\mathbf{u} = [P_1, P_2 \ldots, P_n]^\top \tag{16}$$

The disturbance vector (i.e. non controllable input) d includes solar radiation and external temperature:

$$\mathbf{d} = [G, T_{\text{ext}}]^\top \tag{17}$$

The output vector y consists of the temperatures in each zone:

$$\mathbf{y} = [T_1, T_2, \ldots, T_n]^\top \tag{18}$$

The physical parameter vector $\boldsymbol{\theta}_{phy}$ includes capacitance, resistance and equivalent surfaces:

$$\boldsymbol{\theta}_{phy} = [R_{1,2}, ..., R_{n,n-1}, R_{1,ex}, ..., R_{n,ex}, C_1, ..., C_n, S_1, ..., S_n]^\top \tag{19}$$

The continuous state-space model is then expressed as:

$$\dot{\mathbf{y}} = f_{RC}(\boldsymbol{\theta}_{phy}, \mathbf{y}, \mathbf{u}, \mathbf{d}) = \mathbf{A}(\boldsymbol{\theta}_{phy})\mathbf{y} + \mathbf{B}_1(\boldsymbol{\theta}_{phy})\mathbf{u} + \mathbf{B}_2(\boldsymbol{\theta}_{phy})\mathbf{d} \tag{20a}$$

$$\mathbf{A}(\boldsymbol{\theta}_{phy}) = \begin{bmatrix} -(\frac{1}{R_{1,ex}C_1} + \Sigma_{i\in N_{ad,1}} \frac{1}{R_{1,i}C_1}) & \frac{Ad_{1,2}}{R_{1,2}C_1} & \cdots & \frac{Ad_{1,n}}{R_{1,n}C_1} \\ \frac{Ad_{2,1}}{R_{2,1}C_2} & -(\frac{1}{R_{2,ex}C_2} + \Sigma_{i\in N_{ad,2}} \frac{1}{R_{2,i}C_2}) & \cdots & \frac{Ad_{2,n}}{R_{2,n}C_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{Ad_{n,1}}{R_{n,1}C_n} & \frac{Ad_{n,2}}{R_{n,2}C_n} & \cdots & -(\frac{1}{R_{n,ex}C_n} + \Sigma_{i\in N_{ad,n}} \frac{1}{R_{n,i}C_n}) \end{bmatrix} \tag{20b}$$

$$\mathbf{B}_1(\boldsymbol{\theta}_{phy}) = \begin{bmatrix} \frac{1}{C_1} & 0 & \cdots & 0 \\ 0 & \frac{1}{C_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \frac{1}{C_n} \end{bmatrix}, \qquad \mathbf{B_2}(\boldsymbol{\theta}_{\mathbf{phy}}) = \begin{bmatrix} \frac{S_1}{C_1} & \frac{1}{R_{1,ex}C_1} \\ \frac{S_2}{C_2} & \frac{1}{R_{2,ex}C_2} \\ \vdots & \vdots \\ \frac{S_n}{C_n} & \frac{1}{R_{n,ex}C_n} \end{bmatrix} \tag{20c}$$

where $Ad_{i,j}$ represents the adjacency of two rooms, it equals 1 if connected, 0 if not. $\mathbf{A}(\boldsymbol{\theta}_{phy})$ represents the system dynamics matrix, $\mathbf{B_1}(\boldsymbol{\theta}_{phy})$ is the input matrix, and $\mathbf{B_2}(\boldsymbol{\theta}_{phy})$ is the disturbance matrix.

Discretization and identification objective.

Based on the ODE in equation (20a), a discretization step is needed to perform parameter estimation over a finite data window. We apply the forward Euler method to approximate the time derivative:

$$\dot{\mathbf{y}}[k] \approx \frac{\mathbf{y}[k+1] - \mathbf{y}[k]}{\Delta t} \tag{21}$$

where Δt is the sampling period. Although the Euler method is not the most accurate numerical integration scheme, it is preferred here over higher-order methods such as fourth-order Runge-Kutta. The reasons are twofold:

1. Euler discretization yields a model that is linear with respect to a reparameterized set of variables, enabling the identification problem to be cast as a convex quadratic program;

2. At the same time, the use of Euler discretization does not introduce significant loss of accuracy in the context of building thermal dynamics, given the typical sampling rates.

This property is also beneficial for MPC, where the same discretized model is embedded in the optimization, significantly reducing the computational cost.

The parameter estimation objective minimizes the prediction error in a least-squares sense:

$$\boldsymbol{\theta}^*_{phy} = \arg\min_{\boldsymbol{\theta}_{phy}} \sum_{k=1}^{N} \left\| \frac{\mathbf{y}[k+1] - \mathbf{y}[k]}{\Delta t} - f_{RC}(\boldsymbol{\theta}_{phy},\, \mathbf{y}[k],\, \mathbf{u}[k],\, \mathbf{d}[k]) \right\|_2^2 \tag{22}$$

However, the physical parameter vector $\boldsymbol{\theta}_{phy}$ enters the state-space matrices through products such as $\frac{1}{R_{i,j}C_i}$, making the model nonlinear in $\boldsymbol{\theta}_{phy}$. To obtain a linear least-squares formulation, we

introduce an intermediate parameter vector $\boldsymbol{\theta}_{opt}$ that groups these products. For zone i, this vector is defined as:

$$\theta_{opt,i} = \left[-\left(\frac{1}{R_{i,ex}C_i} + \sum_{j \in N_{ad,i}} \frac{1}{R_{i,j}C_i} \right),\ \frac{1}{R_{i,j_1}C_i},\ \cdots,\ \frac{1}{R_{i,j_n}C_i},\ \frac{1}{C_i},\ \frac{S_i}{C_i},\ \frac{1}{R_{i,ex}C_i} \right]^\top \tag{23}$$

With this reparameterization, the model output for zone i becomes linear in $\theta_{opt,i}$. For zone i, the output vector and data matrix are:

$$Y_i = \begin{bmatrix} T_i[2] - T_i[1] \\ T_i[3] - T_i[2] \\ \vdots \\ T_i[N] - T_i[N-1] \end{bmatrix} \tag{24}$$

$$X_i = \Delta t \cdot \begin{bmatrix} T_i[1] & T_{i,j_1}[1] & \cdots & T_{i,j_n}[1] & P_i[1] & G[1] & T_{ext}[1] \\ T_i[2] & T_{i,j_1}[2] & \cdots & T_{i,j_n}[2] & P_i[2] & G[2] & T_{ext}[2] \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ T_i[N-1] & T_{i,j_1}[N-1] & \cdots & T_{i,j_n}[N-1] & P_i[N-1] & G[N-1] & T_{ext}[N-1] \end{bmatrix} \tag{25}$$

Note that the columns of X_i contain measurements from all adjacent zones $j \in N_{ad,i}$, reflecting the centralized use of multi-zone data even though the parameter vector is structured per zone. To estimate all zone parameters simultaneously, the multi-zone least-squares problem is assembled by stacking the zone-level vectors and matrices:

$$\mathbf{Y} = \begin{bmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{bmatrix}, \qquad \mathbf{X} = \begin{bmatrix} X_1 & 0 & \cdots & 0 \\ 0 & X_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & X_n \end{bmatrix}, \qquad \boldsymbol{\theta}_{opt} = \begin{bmatrix} \theta_{opt,1} \\ \theta_{opt,2} \\ \vdots \\ \theta_{opt,n} \end{bmatrix} \tag{26}$$

With definition above, the optimization problem of multi-zone 1st-order RC model identification is:

Optimization Problem II.1 Multi-zone 1st-order RC model parameter identification

$$\boldsymbol{\theta}^*_{opt} = \arg\min_{\boldsymbol{\theta}_{opt}} \quad \|\mathbf{Y} - \mathbf{X}\,\boldsymbol{\theta}_{opt}\|_2^2 \tag{27a}$$

$$\text{s.t.} \quad \forall i, j \in N_{all},$$

$$\theta_{opt,i,1} < 0, \tag{27b}$$

$$\theta_{opt,i,\,2:end} > 0, \tag{27c}$$

$$\theta_{opt,i,1} + \sum_{j=2}^{n+1} \theta_{opt,i,j} + \theta_{opt,i,n+3} = 0, \tag{27d}$$

$$\theta_{opt,i,j} = 0, \quad \forall j \notin N_{ad,i}. \tag{27e}$$

The first two constraints ensure that thermal resistances and capacitances remain positive. Constraint (27d) enforces energy conservation at each node: the self-coupling term equals the sum of all outgoing

couplings. The last constraint enforces the known building topology by setting to zero the coupling terms between non-adjacent zones.

Once $\boldsymbol{\theta}^*_{opt}$ is obtained, the physical parameters $\boldsymbol{\theta}_{phy}$ can be recovered through the following bijection. Let m denote the number of adjacent zones of zone i, and index the entries of $\theta_{opt,i}$ from 1:

$$R_{i,ex} = \theta_{opt,i,n+1} \Big/ \theta_{opt,i,n+3} \tag{28a}$$

$$R_{i,j} = \theta_{opt,i,n+1} \Big/ \theta_{opt,i,j+1}, \quad j \in N_{ad,i} \tag{28b}$$

$$C_i = 1 \Big/ \theta_{opt,i,n+1} \tag{28c}$$

$$S_i = \theta_{opt,i,n+2} \Big/ \theta_{opt,i,n+1} \tag{28d}$$

Optimization Problem II.1 is a Convex Quadratic Program (QP) with linear equality and inequality constraints. As in [34], this constrained linear least-squares problem can be solved using standard QP solvers, such as the `lsqlin()` function in the Matlab Optimization Toolbox [92]. The objective function is convex and all constraints are linear, which guarantees convergence to the global optimum. This simplicity—low dimension and convexity—is an important asset for scalable deployment in real-world applications.

The identification procedure presented above applies to the first-order RC model, where all states (zone air temperatures) are directly measured. For higher-order models (second or third order), the state vector includes unmeasured quantities such as wall temperatures $T_{wex,i}$. 2nd order RC model identification is showcased in appendix A. This creates a joint estimation problem where both states and parameters must be identified simultaneously. The difficulty is that terms such as $\frac{1}{R_{ex}} T_{wex}$ involve products of unknown parameters and unknown states, making the problem nonlinear. One alternative is to estimate unmeasured temperatures using observers such as the extended Kalman filter. However, the estimated states are projections onto a reduced-order model and may deviate from physical reality, especially under limited data conditions. This deviation weakens the physical interpretability of the gray-box model. For this reason, the first-order model with its direct linear formulation remains the preferred choice for the identification framework used in this work.

Rearranging the Euler discretization formula, the discrete-time prediction model for MPC is obtained:

$$\mathbf{y}[k+1] = \Delta t \cdot f_{RC}(\boldsymbol{\theta}_{phy},\, \mathbf{y}[k],\, \mathbf{u}[k],\, \hat{\mathbf{d}}[k]) + \mathbf{y}[k] = f_{d,RC}(\boldsymbol{\theta}_{phy},\, \mathbf{y}[k],\, \mathbf{u}[k],\, \hat{\mathbf{d}}[k]) \tag{29}$$

where f_d denotes the discrete-time model used for prediction in MPC. This formulation directly predicts the zone temperatures at the next time step $\mathbf{y}[k+1]$ given the current state $\mathbf{y}[k]$, control inputs $\mathbf{u}[k]$, and $\hat{\mathbf{d}}[k]$ represents forecast disturbances including external temperature and solar radiation. The sampling period Δt is typically the supervisory control time step (15–30 minutes for building thermal systems). This first order multi-zone RC (1RCMZ) model is the prediction component embedded in the MPC.

Subspace State Space Identification

Black-box models learn system dynamics directly from measured data without incorporating physical assumptions or prior structural knowledge. Among black-box approaches, subspace identification methods is recognized as a robust tool for deriving state-space representations directly

from input-output data[93]. Unlike autoregressive-based models, subspace methods offer superior approximation performance and require reduced modeling effort [16, 17]. The 4SID method, specifically implemented through the N4SID algorithm, provides a computationally efficient procedure for identifying linear time-invariant systems while maintaining the same level of mathematical complexity as gray-box models with equivalent system order. This equivalence in complexity ensures that both modeling paradigms impose comparable computational burdens during online optimization within the MPC framework.

Subspace identification has been applied to building control across various contexts. Early work demonstrated its suitability as a prediction model within MPC frameworks for building climate control [94]. Subspace models derived directly from building energy simulation tools were shown to enable MPC design without requiring real measurement campaigns [17]. Beyond temperature regulation, subspace identification has also proven effective for heat load prediction in heating buildings [84].

The N4SID algorithm operates on Hankel matrices constructed from input-output measurement sequences. A singular value decomposition (SVD) of measured data matrix reveals the system order n as the number of significant singular values, and provides estimates of the extended observability matrix and the state sequence. From these estimates, the discrete-time state-space matrices (A_d, B_d, C_d, D_d) are recovered using the shift structure of the observability matrix and least-squares regression.

The resulting identified model takes the form of a one-step-ahead predictor:

$$\mathbf{y}[k+1] = f_{d,4SID}(\mathbf{y}[k],\, \mathbf{u}[k],\, \hat{\mathbf{d}}[k]) \tag{30}$$

where $f_{d,\text{4SID}}$ denotes the identified discrete-time state-space model. This formulation enables direct integration into MPC as a prediction constraint, where the multi-step prediction horizon is handled by recursive application within the receding horizon framework.

2.2 Practical Implementation Considerations

The practical deployment of 4SID for building thermal modeling requires attention to several implementation aspects that influence identification quality and control performance. Matlab's System Identification Toolbox provides the `n4sid()` function, which implements efficient singular value decomposition and numerical optimization procedures [30]. The Python-based SIPPY (System Identification Package for Python) library offers an open-source alternative implementation of N4SID algorithms [95]. This computational infrastructure reduces the implementation burden compared to manual algorithm implementation.

Model order selection is a critical decision in 4SID application. Higher-order models provide enhanced approximation capabilities for capturing complex thermal dynamics, but increase susceptibility to overfitting when training data are limited or exhibit insufficient excitation. The model order determines the dimension of the identified state-space representation and should be selected through validation procedures. Cross-validation techniques, information-theoretic criteria such as Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC), and singular value analysis of the Hankel matrix provide approaches for model order determination. In building thermal management with limited historical data, conservative model order selection often yields more robust predictive performance than aggressive high-order fitting.

Data preprocessing impacts identification quality. Proper detrending removes low-frequency drift that can bias parameter estimates, while filtering eliminates high-frequency measurement noise

that corrupts the subspace projections. Sampling time selection must balance capturing relevant thermal dynamics against numerical conditioning of the Hankel matrices. For building applications, sampling intervals ranging from 5 to 15 minutes typically provide adequate temporal resolution while maintaining computational tractability.

Training dataset length and quality represent limiting factors in 4SID performance. As demonstrated in validation studies, 4SID models often require longer training periods (approximately 10 days) compared to gray-box alternatives (1 day) to achieve comparable prediction accuracy [86]. This extended data requirement reflects the purely empirical nature of black-box identification, which must learn both system structure and parameters simultaneously from measurements. Additionally, 4SID exhibits greater sensitivity to dataset selection and operating conditions, occasionally producing outlier performance with certain training configurations even when overall data length appears sufficient. This variability necessitates robust validation protocols and potentially multiple identification trials with different data segments to ensure reliable model quality.

Building on the bias-variance tradeoff framework discussed earlier, gray-box RC and black-box 4SID models represent different positions in this tradeoff. RC imposes physical structure through energy conservation, reducing variance at the cost of potential bias from model assumptions. 4SID provides greater flexibility to capture complex dynamics, reducing bias but increasing variance sensitivity to training data. This chapter will investigate how these models perform in practice under limited operational data.

3 Hierarchical MPC Design

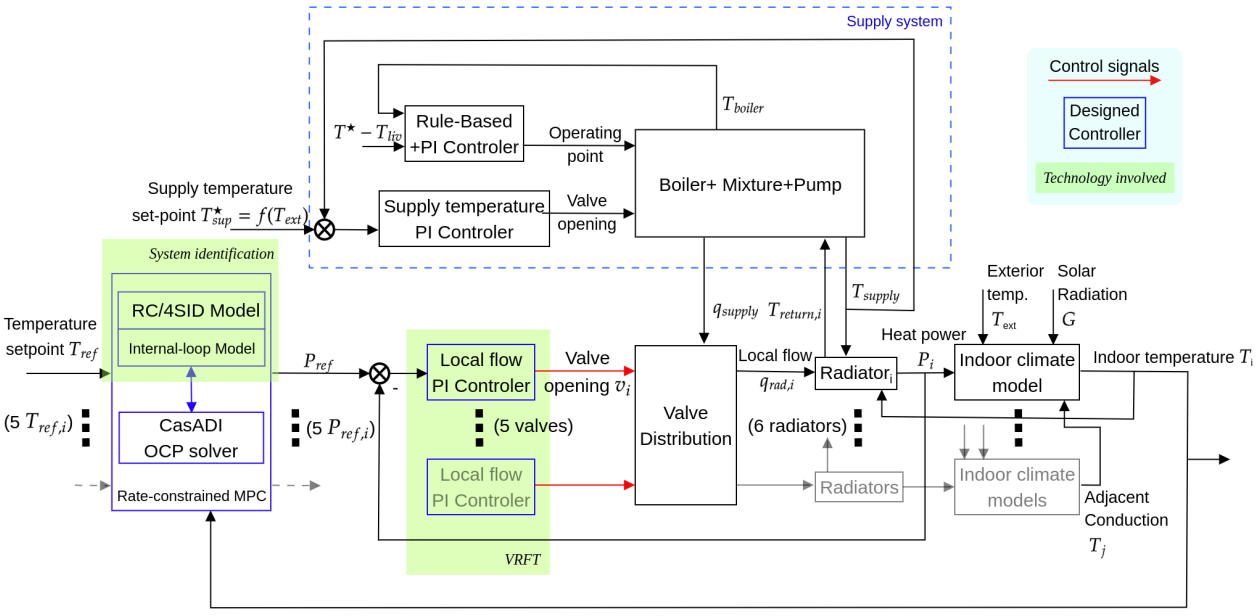

Figure II.4: The proposed hierarchical control scheme applied to MZRH test case.

This section uses the MZRH test case (see section I.2.3) to illustrate the proposed control scheme. As shown in the figure II.4, the system consists of two nested closed-loop layers. The inner layer tracks thermal power across five zones. The five radiators are coupled through water flow distribution within

the hot water circulation system and decoupled via five independent PI controllers. Considering that power reduction processes may cause control outputs to saturate at the lower limit for long periods, these PI controllers incorporate anti-windup design through VRFT. This inner loop also mitigates disturbances from the boiler system, such as fluctuations in inlet water temperature and pressure. The reference power signals for the five PI controllers are optimized by MPC, balancing precise tracking of the reference room temperature with energy consumption. The multi-zone indoor climate model is obtained via RC identification or subspace identification. To incorporate radiator residual heat into optimization and fully take advantage of building thermal inertia, rate constraints are added to the MPC. This design enhances energy efficiency while preserving the quadratic nature of the MPC optimization problem, leading to higher solution efficiency.

3.1 Supervisory-level MPC Formulation

Challenge for supervisory MPC in building thermal control

Assuming an ideal tracking level control, the typical optimization solved by MPC at this level is

---

Optimization Problem II.2 Nominal MPC

$$\min_{\mathbf{u}_{ref}} \quad \sum_{j=0}^{N-1} \left(\mathbf{e}[k+j+1]^\top Q\, \mathbf{e}[k+j+1] + \mathbf{u}_{ref}[k+j]^\top R\, \mathbf{u}_{ref}[k+j]\right) \tag{31a}$$

$$\text{s.t.} \quad \mathbf{e}[k+j+1] = \mathbf{y}_{ref}[k+j+1] - \hat{\mathbf{y}}[k+j+1|k] \tag{31b}$$

$$\hat{\mathbf{y}}[k+j+1|k+j] = f_d(\mathbf{y}[k+j], \mathbf{u}_{ref}[k+j], \hat{\mathbf{d}}[k+j]) \tag{31c}$$

$$\underline{\mathbf{u}} \leq \mathbf{u}_{ref}[k+j] \leq \bar{\mathbf{u}}, \quad \forall j \in \{0, 1, \ldots, N-1\} \tag{31d}$$

At each discrete time step k, the MPC solves Optimization Problem II.2. It searches for the optimal power reference trajectory $\mathbf{u}_{ref}$ within the lower bounds $\underline{\mathbf{u}}$ and upper bounds $\bar{\mathbf{u}}$. The first term of the quadratic cost function (31a) quantifies thermal comfort through the weighted norm of the temperature error $\mathbf{e}$ over the prediction horizon of N steps. The second term penalizes the requested thermal power. Q and R are positive definite diagonal weight matrices. As the thermal requirements are common to each zone, Q is set to the identity matrix I and R is defined as wI, where w is a scalar weighting parameter. $\hat{\mathbf{y}}[k+j+1|k+j]$ denotes the one-step-ahead predicted output at step $k+j+1$, computed from the state at step $k+j$ using the indoor climate model f_d. The predicted trajectory is obtained by recursive application of f_d over the horizon, starting from the current measurement $\mathbf{y}[k]$. $\hat{\mathbf{d}}[k+j]$ is the forecast disturbance, namely solar radiation and exterior temperature. This formulation is referred to as 'Nominal MPC' in this chapter.

Nominal MPC provides a general mathematical framework for control problems, but assumes perfect tracking of power references at the local level - an assumption that may not hold in practice. Figure II.5 illustrates the thermal power response characteristics of the internal control loop. The graph demonstrates that both the blue and red thermal power reference trajectories can result in the same actual power response, represented by the black curve. In the nominal MPC approach, the controller would request an abrupt change in power demand (red dashed curve) and then wait for the

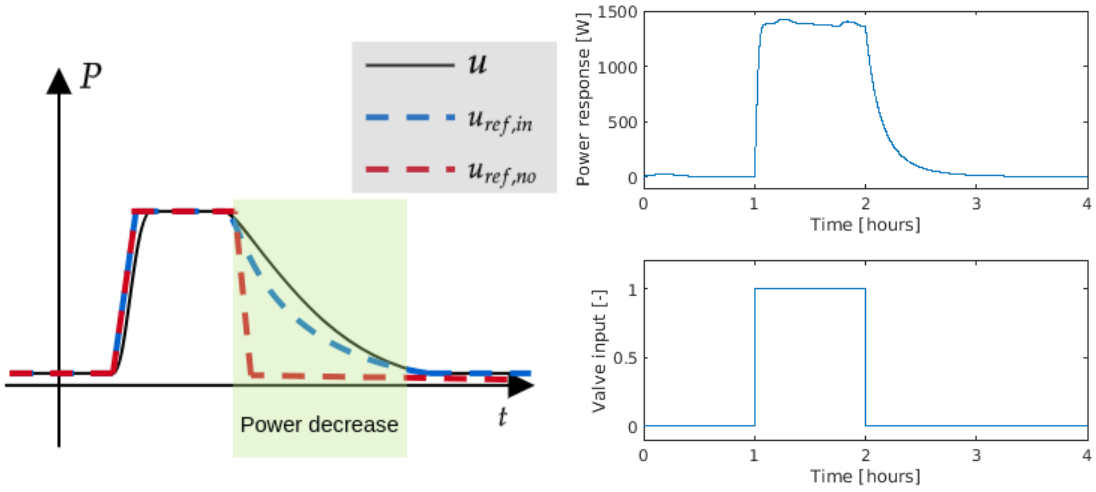

Figure II.5: Power response of heating equipment. A nominal MPC suggests power reference $u_{ref,no}$ that waits for power decrease process. An internal loop model contributes to a closer prediction to local prediction, like $u_{ref,in}$. Right side shows the MZRH test case radiator power response as an example. Top: Power output of radiator. Bottom: valve input.

actual power to decrease. This behavior stems from the operating mechanism of the equipment. For instance, water-based HVAC exhibit thermal inertia that can result in a gradual power fading process due to the indoor convective heat transfer mechanisms, potentially spanning several MPC time steps. In contrast, an MPC with a input rate constraint (blue dashed curve) can fully account for the internal loop dynamics [96].

Identification for Power Dynamics Constraint

By modeling and constraining the gradual thermal power evolution, the MPC can react to set-point changes earlier, thereby optimizing the utilization of residual energy in the system and improving overall energy efficiency. An internal loop model is proposed as a rate constraint:

$$\mathbf{u}_{ref}[k+1] - \mathbf{u}_{ref}[k] \geq A_{in}\mathbf{u}_{ref}[k] - \mathbf{u}_{ref}[k] \tag{32a}$$

$$\Rightarrow \mathbf{u}_{ref}[k+1] - A_{in}\mathbf{u}_{ref}[k] = \mathbf{a} \geq 0 \tag{32b}$$

where A_{in} is a stable diagonal matrix to be estimated (diagonal elements smaller than 1), representing the thermal inertia of the heating device. This model approximates the internal loops as discrete first-order systems. Equation (32b) is incorporated into the optimization as a dynamic constraint. It establishes a lower bound on the rate of decrease for the reference power vector $\mathbf{u}_{ref}$. The auxiliary variable $\mathbf{a}$ allows for power increases, acting as a non-negative slack variable. When $\mathbf{a}[k] = 0$, $\mathbf{u}_{ref}[k+1]$ reaches its minimum value $A_{in}\mathbf{u}_{ref}[k]$. $\mathbf{a}$ is not upper-bounded, reflecting the rapid power increase capability of the system. Consequently, the optimization always permits $\mathbf{u}_{ref}[k+1]$ to exceed $\mathbf{u}_{ref}[k]$, subject to the maximum power limit of the device. A_{in}^* is obtained by solving the following least-squares problem:

Optimization Problem II.3 Tracking Loop Power Decrease Dynamic Model Identification

$$\min_{A_{in}} \quad \sum_{k \in D_{in}} \|\mathbf{u}[k+1] - A_{in}\mathbf{u}[k]\|^2 \tag{33a}$$

$$\text{with} \quad D_{in} = \{k \mid \mathbf{u}[k] - \mathbf{u}[k+1] \geq 0, \quad k \in \{1, \ldots, n_{data}\}\} \tag{33b}$$

Only data relating to power reduction is required to identify A_{in}. The training set index is defined in equation (33b).

Complete Rate-Constrained MPC

The complete rate-constrained MPC is then formulated as:

Optimization Problem II.4 Rate-constrained MPC

$$\min_{\mathbf{u}_{ref}} \quad \sum_{j=0}^{N-1} \left(\mathbf{e}[k+j+1]^\top Q\, \mathbf{e}[k+j+1] + \mathbf{u}_{ref}[k+j]^\top R\, \mathbf{u}_{ref}[k+j] \right) \tag{34a}$$

$$\text{s.t.} \quad \mathbf{e}[k+j+1] = \mathbf{y}_{ref}[k+j+1] - \hat{\mathbf{y}}[k+j+1|k+j] \tag{34b}$$

$$\hat{\mathbf{y}}[k+j+1|k+j] = f_d(\mathbf{y}[k+j], \mathbf{u}_{ref}[k+j], \hat{\mathbf{d}}[k+j]) \tag{34c}$$

$$\mathbf{u}_{ref}[k+j+1] = A_{in}^* \mathbf{u}_{ref}[k+j] + \mathbf{a} \tag{34d}$$

$$\mathbf{a} \geq 0 \tag{34e}$$

$$\underline{\mathbf{u}} \leq \mathbf{u}_{ref}[k+j] \leq \bar{\mathbf{u}}, \quad \forall j \in \{0, 1, \ldots, N-1\} \tag{34f}$$

where Optimization Problem 34d and Optimization Problem 34e encode the internal loop dynamics from Optimization Problem 32b. This optimization problem is solved using CasADi [97], which enables efficient implementation through automatic differentiation and mathematical-like syntax for problem definition. The solution employs multiple shooting to enhance both feasibility and computational speed.

An alternative approach to account for asymmetric power dynamics is to add an input increment penalty $\|\Delta\mathbf{u}_{ref}[k+j]\|^2$ to the cost function, where $\Delta\mathbf{u}_{ref}[k+j] = \mathbf{u}_{ref}[k+j+1] - \mathbf{u}_{ref}[k+j]$. This penalizes large variations in the power reference. However, in HVAC systems, the power dynamics are asymmetric: power can increase rapidly (by opening a valve), but can only decrease slowly due to thermal inertia in the heating circuit. A symmetric quadratic penalty on $\Delta\mathbf{u}_{ref}$ cannot capture this asymmetry. To do so, the penalty would need to apply different weights depending on the sign of the increment, which requires conditional expressions of the form:

$$J_{\Delta u} = \sum_{j=0}^{N-2} \begin{cases} w^+ \|\Delta\mathbf{u}_{ref}[k+j]\|^2 & \text{if } \Delta\mathbf{u}_{ref}[k+j] \geq 0 \\ w^- \|\Delta\mathbf{u}_{ref}[k+j]\|^2 & \text{if } \Delta\mathbf{u}_{ref}[k+j] < 0 \end{cases} \tag{35}$$

Such conditional cost introduces non-smoothness into the objective function, making the optimization non-convex and harder to solve. In contrast, the proposed rate constraint in Optimization Problem II.4

handles the asymmetry through a linear inequality constraint, preserving the convex quadratic structure of the optimization.

3.2 Tracking-level Controller Tuning via VRFT

We consider a multivariable power tracking control problem arising in hydronic HVAC systems equipped with multiple radiators. Due to hydraulic and thermal interactions within the HVAC network, the delivered powers are coupled, even when the controllers are implemented in a decentralized manner. The control objective is the regulation of radiator power, rather than direct temperature control. Each radiator tracks a prescribed power reference while ensuring smooth actuation, bounded control effort, and robustness with respect to inter-radiator coupling. This power tracking layer operates as an inner control loop, decoupled from higher-level thermal supervision or energy management strategies.

To address this problem, a multivariable Virtual Reference Feedback Tuning (VRFT) framework is adopted. VRFT enables the direct tuning of controllers from experimental data without explicit identification of the HVAC plant model. This feature is particularly suited for HVAC applications, where first-principles modeling of hydraulic couplings and actuator dynamics is often complex, uncertain, and time-varying. Considering that actuators always have physical saturation, [98] developed VRFT tuning for anti-windup PID controllers. This data-driven approach supports diverse control schemes by designing controller arrays and reference models.

A diagonal reference model M_{ref} is selected to specify the desired closed-loop behavior of each power channel. Each diagonal element is defined as a stable, minimum-phase first-order transfer function augmented with a fast zero, whose time constant is chosen one order of magnitude smaller than the dominant pole. This design preserves the invertibility and numerical conditioning required by VRFT, while approximating integral-like behavior at low frequencies. As a result, steady-state power tracking errors induced by structural biases or slow variations in the HVAC network can be mitigated without introducing marginally stable dynamics in the reference model.

Although the reference model is diagonal, the VRFT tuning is performed in a multivariable setting, using measured power data from all radiators simultaneously. The learned controllers implicitly account for the coupling effects induced by the HVAC system, despite the decentralized implementation of the control law. This represents a deliberate trade-off between performance and complexity, whereby local closed-loop behavior is explicitly shaped, while inter-channel interactions are handled implicitly through robustness.

The controller structure is chosen as a discrete-time PID with anti-windup compensation, parameterized linearly with respect to the tuning variables. This choice ensures compatibility with the VRFT framework, facilitates reliable parameter estimation from data, and reflects common industrial practice in HVAC control. The inclusion of anti-windup action is essential to prevent biased tuning in the presence of actuator saturation, which is unavoidable in power-limited heating systems.

The tracking-level control problem is defined as follows:

- The HVAC system G is a multiple-input multiple-output system with an input vector $\mathbf{v} = [v_1, \ldots, v_{n_z}]^\top$ (valve openings) and an output vector $\mathbf{P} = [P_1, \ldots, P_{n_z}]^\top$ (thermal powers).

- A set of discrete controllers C achieves the power tracking goal. The controller output is given by equation (36):

$$\begin{aligned}\mathbf{v}[k] &= C_{PID}(q^{-1})\boldsymbol{\theta_{PID}}\,\mathbf{e}_P[k] + C_{aw}(q^{-1})\boldsymbol{\theta_{aw}}\,\Delta\mathbf{v}_{sat}[k] \\ &= \left[1,\ \frac{T_{PID}}{2}\frac{1+q^{-1}}{1-q^{-1}},\ \frac{2}{T_{PID}}\frac{1-q^{-1}}{3-q^{-1}}\right]\boldsymbol{\theta_{PID}}\,\mathbf{e}_P[k] \\ &\quad + \left(\frac{T_{PID}}{2}\frac{1+q^{-1}}{1-q^{-1}}\right)\boldsymbol{\theta_{aw}}\,\Delta\mathbf{v}_{sat}[k]\end{aligned} \tag{36}$$

  where q denotes the shift operator and T_{PID} is the PID discrete time step. $\mathbf{e}_P = [e_{P,1}, \ldots, e_{P,n_z}]^\top$ is the power tracking error vector. C_{PID} is the discrete transfer function vector of the PID controller. The parameter vector of the PID controller array is $\boldsymbol{\theta_{PID}} = [\theta_{PID,1}, \ldots, \theta_{PID,n_z}]^\top$, where $\theta_{PID,i} = [K_{P,i}, K_{I,i}, K_{D,i}]^\top$ are the proportional, integral, and derivative gains for zone i. C_{aw} is the transfer function of the anti-windup branch with gains $\boldsymbol{\theta_{aw}} = [K_{aw,1}, \ldots, K_{aw,n_z}]^\top$. $\Delta\mathbf{v}_{sat}$ is the difference between the controller output before and after saturation.

- The diagonal reference model M_{ref} specifies the desired closed-loop response. $M_{ref}^\dagger$ denotes its left inverse, satisfying $M_{ref}^\dagger M_{ref} = I$.

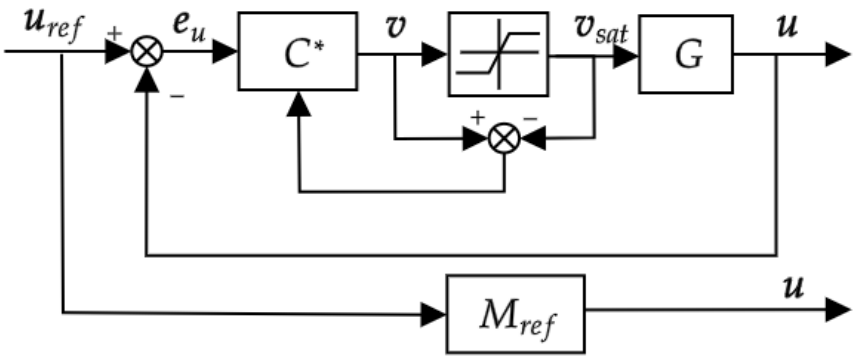

Figure II.6: Block diagram of VRFT process. C^* represents an ideal controller that drives the closed-loop to achieve the desired performance M_{ref}.

Figure II.6 illustrates the VRFT method. This approach optimizes a given controller class (e.g., PID with anti-windup) to approximate an ideal controller C^* with respect to the reference model using only input-output data from the system G. The ideal controller C^* is defined as the controller that, if inserted in the closed-loop during the data collection experiment, would have produced exactly the desired behavior. VRFT generates a virtual reference $\mathbf{P}_{ref,vr}$ by reverse-simulating the desired closed-loop response through $M_{ref}^\dagger$. This virtual reference allows VRFT to construct the controller input data without requiring an explicit system model.

The optimal parameters $\boldsymbol{\theta_{PID}}$ and $\boldsymbol{\theta_{aw}}$ are obtained by solving:

Optimization Problem II.5 VRFT Parameter Estimation

$$\min_{\boldsymbol{\theta_{PID}}, \boldsymbol{\theta_{aw}}} \quad \sum_{k=1}^{n_{id}} \|\boldsymbol{\varepsilon}[k]\|^2 \tag{37a}$$

$$\text{s.t.} \quad \boldsymbol{\varepsilon}[k] = \mathbf{v}[k] - C_{PID}(q^{-1})\boldsymbol{\theta_{PID}}\,\mathbf{e}_P[k] - C_{aw}(q^{-1})\boldsymbol{\theta_{aw}}\,\Delta\mathbf{v}_{sat}[k] \tag{37b}$$

$$\mathbf{P}[k] = M_{ref}(q^{-1})\mathbf{P}_{ref,vr}[k] \tag{37c}$$

$$\mathbf{e}_P[k] = \mathbf{P}_{ref,vr}[k] - \mathbf{P}[k] \tag{37d}$$

The tracking error in Optimization Problem (37d) can be expressed as:

$$\mathbf{e}_P[k] = (M_{ref}^{\dagger}(q^{-1}) - I)\,\mathbf{P}[k] \tag{38}$$

When generating the dataset in open-loop, inputs should exceed the saturation limits to make $\Delta\mathbf{v}_{sat}$ data informative. The tuning estimates the linear part of the controller, while the saturation is absorbed into the system model. Quadratic programming solvers, such as `quadprog()` in the Matlab Optimization Toolbox, can solve Optimization Problem II.5.

The reference model time constant τ is a critical parameter that balances the response speed with the inherent dynamics of the equipment. It must be fast enough to support the quasi-static power response assumption of the supervisory MPC, yet slow enough to be achievable by the physical system. As an initial guideline, setting τ to $1/10$ of the MPC time step provides a starting point for further refinement. This ratio balances the supervisory level's tracking requirements with actuator longevity, as overly aggressive response times could accelerate equipment wear.

In equation (38), $M_{ref}^{\dagger}$ is used to backward-simulate the virtual reference. If the reference model has more poles than zeros, $M_{ref}^{\dagger}$ becomes non-causal, making simulation infeasible. An approximate solution is to add a zero with a time constant much faster than the poles. A zero 10 times faster is sufficient in practice. In the current case, M_{ref} is taken as a diagonal transfer matrix consisting of several transfer functions

$$M_{ref,i}(q^{-1}) = \text{ZOH}\left(\frac{\frac{1}{10}\tau s + 1}{\tau s + 1}\right) = \frac{\beta_0 + \beta_1 q^{-1}}{1 - q^{-1}e^{-T_{PID}/\tau}} \tag{39}$$

where τ is the desired time constant, T_{PID} is the PID discrete time step, and β_0, β_1 are obtained through zero-order hold classical Zero Order Hold (ZOH) discretization.

This choice is well-suited for VRFT, that tunes the controller using adaptation to experimental data. A fast reference model enables the system to respond more quickly than its natural dynamics while remaining feasible. By keeping the zero sufficiently fast, the risks of overshoot and instability—which could arise if the zero were too slow—are mitigated, thereby enhancing robustness. Finally, the discretized reference model is compatible with the structure of a discrete PID controller, making practical implementation easier

In the proposed structure, local-level variables such as zone temperature, supply water temperature, and supply water flow rate act as disturbances to the tracking-level control. However, these variables are typically outputs of other closed-loop systems within the building, which naturally constrains their variability. Incorporating realistic operating conditions when generating the tuning dataset helps ensure controller effectiveness. The controller parameters should be re-tuned when the operating

point of the tracking control changes. The PI controller structure inherently provides disturbance rejection capabilities suited to building services. This combination of constrained disturbances and robust controller structure results in reliable power regulation.

4 Simulation Validation and Performance Evaluation

The hierarchical control framework and tuning methods presented in the previous sections are generic and applicable to diverse building types and HVAC equipment. The supervisory-level MPC formulation (section II.1) and the VRFT-based tracking controller design (section II.2) do not depend on specific building or equipment characteristics. To validate their generalizability, this scheme is deployed on two distinct BOPTEST test cases introduced in section I.3: the MZRH test case with radiator-based heating, and the MOSA test case with air-handling systems. This validation on architecturally different buildings confirms that the proposed approach achieves equipment-agnostic supervisory control and reliable power tracking across diverse configurations. Results from these validation have been published in [86, 99]. **For the clarity of discussion on RC vs. 4SID model, this result section focus on the MZRH test case.**

The Results section evaluates the proposed hierarchical MPC framework for building thermal management. It is structured in three parts:

1. Validation of indoor climate models (1RCMZ and 4SID).
2. Performance of tracking level controllers tuned by VRFT.
3. Overall performance of the hierarchical control scheme and model comparison in closed-loop.

4.1 Open-loop Validation of Indoor Climate Models

The indoor climate model is trained and validated using baseline control historical data in the test case. The dataset is divided into a training set, covering the period from January 1 to January 21, and a validation set, spanning January 22 to January 28. Both the gray-box (1RCMZ) and black-box models are trained using data regulated by baseline controllers. The dataset sampling time for building multi-zone model identification is $T_{dataset} = 60$ sec. Temperature derivatives $\frac{\mathrm{d}\bar{y}(t_k)}{\mathrm{dt}}$ are approximated by forward Euler differentiation with $\Delta t = 20$ sec. To ensure fairness in comparison, the order of the 4SID model is set to 8, the same as the RC model order. Figure II.7 presents the prediction over a 6-hour horizon for both the 1RCMZ and 4SID models in data validation. The 1RCMZ model demonstrates good approximation during periods of low solar radiation (t=24.8 and t=25). However, its error noticeably increases during strong solar radiation hours (t=24.6 and t=25.6). This highlights a weakness in the 1RCMZ model's handling of solar radiation. The actual radiation process is far more complex than the simple static gain and direct irradiance multiplication used in 1RCMZ, involving time-variant and non-linear factors like sunspot. In contrast to 1RCMZ, the 4SID model shows lower error during high solar radiation periods and higher error when solar radiation is low.

To better evaluate prediction quality, an indicator I_{pred} is defined. This indicator, defined in Equation (40), measures the average Mean Squared Error (MSE) of temperature predictions across all MPC control time steps in the simulation. For each zone i, $I_{\text{pred},i}$ is calculated by first computing the

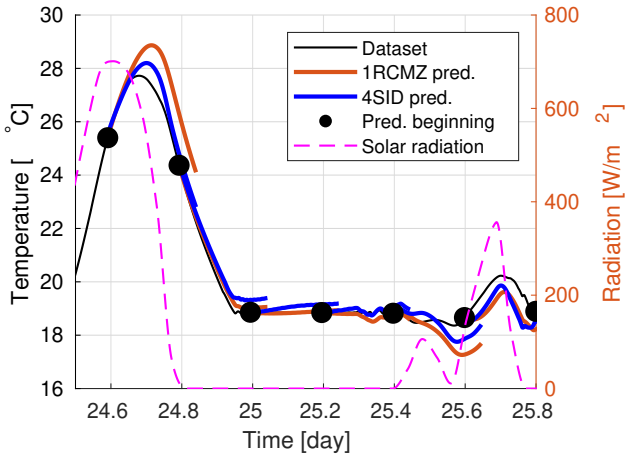

Figure II.7: Data validation of indoor climate models predictions. Temperature of bedroom 1 (Ro1.) is shown as an example. Prediction horizon is 6 hours.

MSE between predicted temperatures $\hat{T}_i$ and actual measured temperatures T_i over the prediction horizon N at each MPC time step. These MSE are then averaged over all MPC time steps in the simulation.

$$I_{\text{pred},i} = \frac{1}{n_{\text{MPC}}} \sum_{k=1}^{n_{\text{MPC}}} \left(\frac{1}{N} \sum_{p=1}^{N} (T_i[k+p] - \hat{T}_i[k+p]))^2 \right) \tag{40}$$

Figure II.8 illustrates the approximation capabilities of the two models, focusing on bedroom 1 (Ro1.), 2 (Ro2.), and 3 (Ro3.) as representative examples. All these three zones are the same in size, but Ro3. is on the opposite side of the house. Therefore, Ro1. and Ro2. share the same radiation condition, Ro3. is different. The figure displays the I_{pred} of three zones predicted by 1RCMZ and 4SID in different forecast horizons.

I_{pred} increases with the length of the forecast horizon for both models, reflecting the accumulation of errors over time. For Ro3. which has different solar radiation conditions, its predicted quality differs from that of Ro1. and Ro2. This confirms that the 1RCMZ modeling error is sensitive to solar radiation. Within the first 4–6 hours, both models exhibit comparable prediction accuracies, ensuring that their similar open-loop performance validates the fairness of the closed-loop comparison.

The sensitivity to training set length for both models is further studied in figure II.9. The results reveal significant differences in the training stage. The 1RCMZ model converges to its best performance within one day's data. In contrast, the 4SID model requires nearly 10 days to converge to a good prediction model. Notably, the MSE for days 1, 2, and 3 are so high they fall outside the graph's range. 1RCMZ demonstrates strong consistency, with its MSE consistently floating between 0.15 and 0.4 across different dataset lengths. The 4SID model also shows more variability. While the overall trend of 4SID's MSE is decreasing with more data, it produces outlier MSE (day 11, 19) with certain training sets. The poor performance on certain datasets can be mitigated by adjusting the order of the 4SID model or modifying the data sampling time. However, these adjustments lead to heavy workload in scalable deployment. These results confirm the 4SID method's higher sensitivity to the sampling configuration compared to 1RCMZ.

4.2 Tracking Controller Performance Assessment

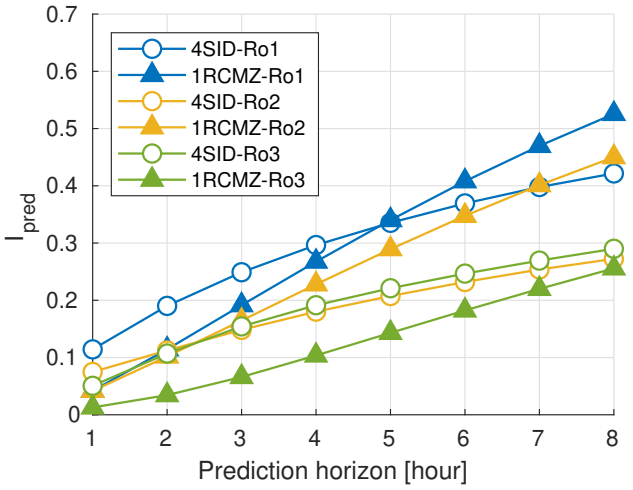

Figure II.8: Error propagation with growth of prediction horizon

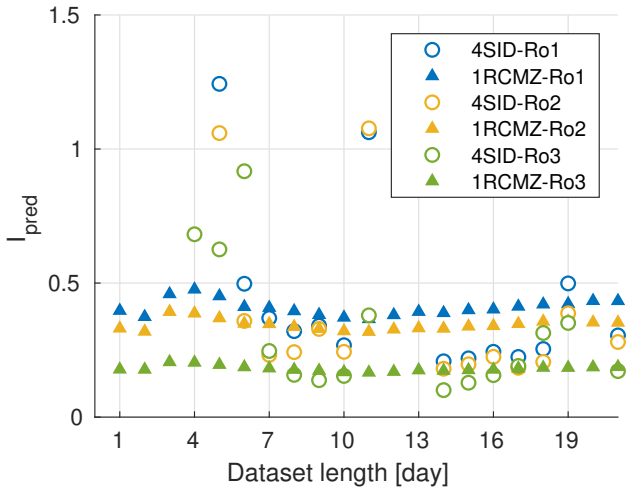

Figure II.9: Identification convergence and training dataset length. In this figure, 6 hour prediction quality is validated in zone Ro1.,Ro2. and Ro3..

Zone	Nominal power	K_P	K_I	K_{aw}
Liv.	1.8 kW	0.0651	$1.9e^{-13}$	$3.4e^{-13}$
Bth.	0.8 kW	2.9579	$3.9e^{-15}$	$8.5e^{-17}$
Ro1.	1 kW	1.9813	$4.1e^{-15}$	$1.1e^{-16}$
Ro2.	0.9 kW	3.0351	$3.2e^{-15}$	$7.2e^{-17}$
Ro3.	1.1 kW	1.9825	$4.0e^{-15}$	$1.1e^{-16}$

Table II.1: VRFT tuned PI controller parameters for each zone. The integral gains K_I and anti-windup gains K_{aw} appear small in magnitude but play a significant role in controller performance. The low numerical values result from the long integration period in the discrete controller implementation and the physical scale of the valve flow dynamics.

In this test case, the designed architecture use PI controllers with anti-windup for power tracking. Table II.1 shows the PI controller parameters for different zones, tuned by VRFT introduced in section II.2 with derivative gain terms removed in optimization. The nominal power indicates radiator size, with the Living room (Liv.) having the largest at 1.8 kW. A clear trend emerges: for larger radiators, K_P is smaller while K_I and K_{aw} are larger. This relationship aligns with PI control principles. A smaller K_P for larger radiators reduces the risk of overshoot, as these radiators have more thermal inertia. The larger K_I compensates for the reduced proportional action, ensuring steady-state errors are eliminated over time. The increased K_{aw} helps prevent integral windup, which is more critical for larger systems with potentially longer saturation periods.

Figure II.10 validates the VRFT-tuned controllers in closed-loop operation. The power response shows fast and accurate tracking during power increase. The valve primarily operates binary-like states (almost fully open or closed) due to the high supply water temperature. Radiator can reach the power reference with a fast water intake and maintain the power with a small water flow. When it maintain a higher power reference (like the fifth square wave), it requires a more obvious valve opening to continuously replenish hot water. Although valve operations in real systems exhibit smoother behavior

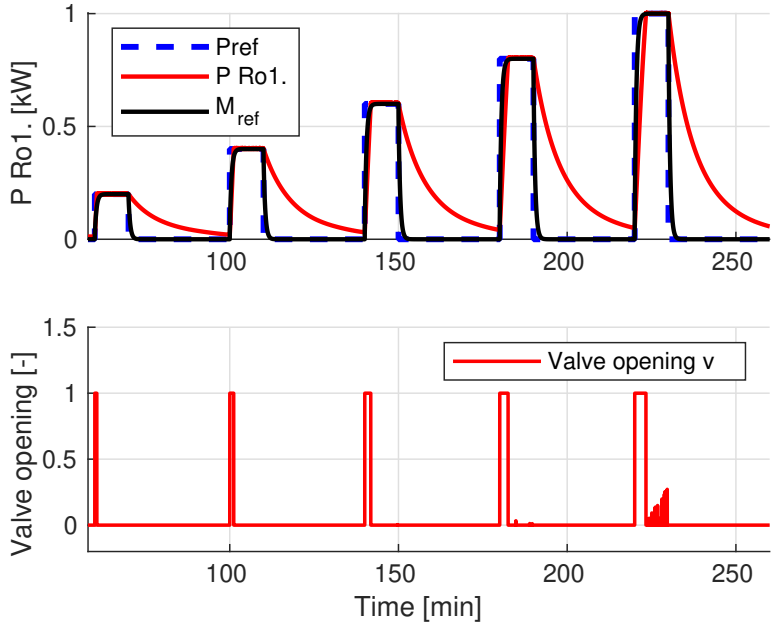

Figure II.10: Internal loop simulation that validates VRFT-tuned controllers. Zone Ro1. is shown as a sample. Hot water temperature, zone temperature and water flow is set to constant. Power reference is square wave with 10 minutes heating and 40 minutes stand-by.

rather than impulse signals shown in the figure, this stand-alone test justify the tracking quality of VRFT-tuned controllers. More complex power dynamics will be present in next subsection.

4.3 Closed-loop Hierarchical Control Performance

In the test case, performance of hierarchical control schemes are further evaluate. In this stage, 1RCMZ and 4SID are integrated into the Rate-constrained MPC formulation in optimization problem II.4. Parameters tuned by VRFT in table II.1 is adopted for local controllers. Parameter setting of Rate-constrained MPC is shown in table II.2.

Parameters	Values	Description
Model	1RCMZ/4SID	Prediction model
T_{MPC}	10 minutes	MPC control time step
H	4/6 hours	MPC prediction horizon
Q	$I_{(8,8)}$	Weighting on state error vector
R	$wI_{(6,6)}$	Weighting on power input vector
Solver	'qpoases'[100]	Solver in MPC
w	$w_0 = 1.6 \times 10^{-4}$	Weighting on objectives

Table II.2: Rate-constrained MPC parameters used in test case

Figure II.11 illustrates the temperature evolution across three zones from January 22 to January 29, using the same climate file, solar radiation, and exterior temperature as the validation dataset. Despite the 4SID model shows a comparable approximation capabilities to 1RCMZ in figure II.8, it has poor

performance in the closed-loop validation. Notably, deviations from the set-point are evident: the Living Room is under-heated, while Room 2 is overheated. When environmental perturbations exceed its extrapolation boundaries, the resulting biased temperature prediction is overestimated/underestimated. In contrast, the RC model demonstrates consistent reliability when integrated into our framework, as its informed knowledge maintain prediction fidelity across wider operating conditions.

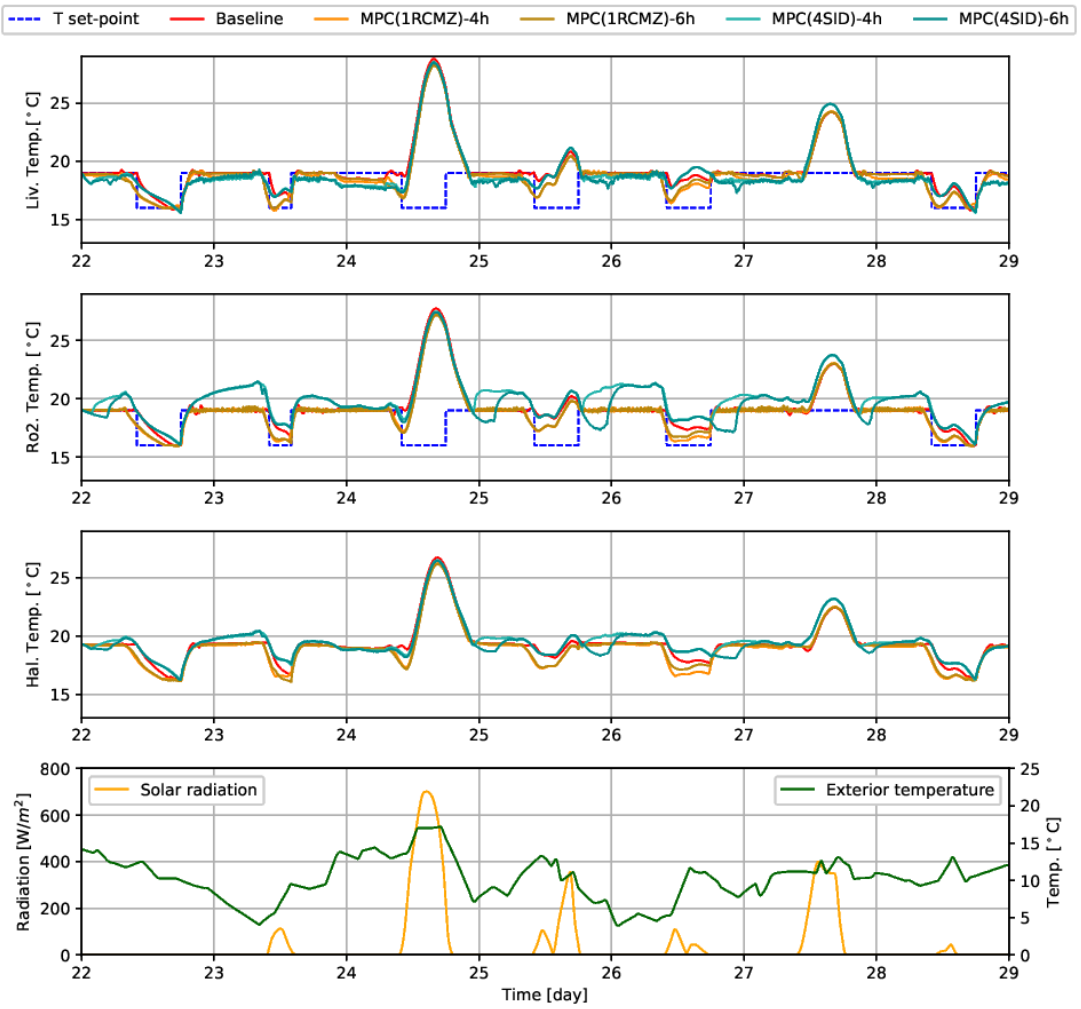

Figure II.11: Comparison of 1RCMZ and 4SID performance in proposed control scheme. The top three figures show the temperatures of the living room (liv.), bedroom (Ro2.), and hallway (hal.). Both models were tested with 4-hour and 6-hour prediction horizons. The bottom figure shows disturbance variation.

Figure II.12 provides a one-day zoom-in to compare the prediction quality of 1RCMZ and 4SID models. Figure II.12a shows that 1RCMZ predicts well during occupied periods, with most predictions aligning with actual temperatures, except at midday. The MPC consistently underestimates indoor temperatures during solar radiation hours, but predictions realign with actual temperatures as solar radiation fades.

In figure II.12b, 4SID predictions deviate significantly from reality. At $t = 24.9$ day, the model overestimates temperature decrease when there's no solar radiation or radiator power. At $t = 25.2$

day, despite Ro2. being overheated, the MPC still demands maximum radiator power $\bar{u} = 800W$, as indicated by u_{ref} Ro2.This inaccurate prediction in Ro2 also affects Liv.. Actual radiator power u Liv. is not enough to achieve the reference u_{ref} Liv. due to excessive demand in Ro2, as Ro2 consumes a disproportionate share of power in the hot water circulation.

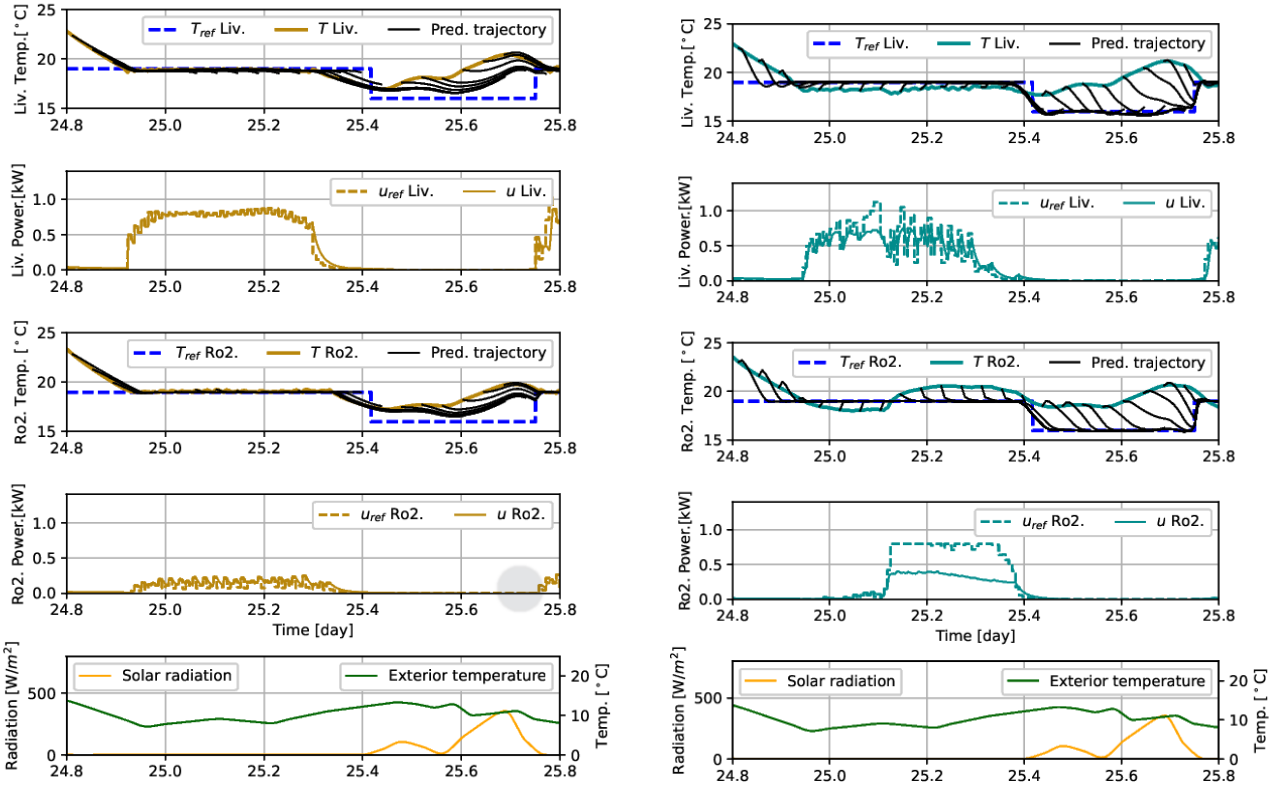

(a) Rate-constrained MPC-1RCMZ

(b) Rate-constrained MPC-4SID

Figure II.12: Zoom-in on the comparison of Rate-constrained MPC formulation with 1RCMZ and 4SID models in test case. Each scheme shows Liv. and Ro2. as samples. Figure exhibits temperature and power responses of each zone and also disturbances of corresponding period. Every hour $(6 * T_{MPC})$ a temperature prediction trajectory is shown as a black line labeled 'Pred. trajectory' in temperature subplots.

Baseline historical data as the training set is very concentrated around the set-point. It is similar to red curve in figure II.11. The model is calibrated over a narrow range. However, the MPC searches for optimal control inputs and state trajectories over a broader range. The optimization process may identify solutions within the 4SID state space that lie outside the reliable region but yield a lower cost in the objective function. Consequently, while the predicted trajectories converge well to the set-point, the power reference generated by the MPC may be inconsistent.

The model prediction quality (evaluated by indicator I_{pred} in equation (40)) in closed-loop is compared in figure II.13. The 4SID model produces higher errors than the 1RCMZ model throughout the simulation. The figure also reveals that the mismatch in Liv. is higher than in other rooms. For 1RCMZ, as explained, the insufficiency lies in the radiation model. The living room has a larger window surface, so it is more sensitive than other zones. For 4SID, predicted trajectories always converge to

the set-point, error is thus larger when temperatures rise in solar radiation hours. Since living room temperature in non-occupied time is always higher than other zones, it has bigger deviations.

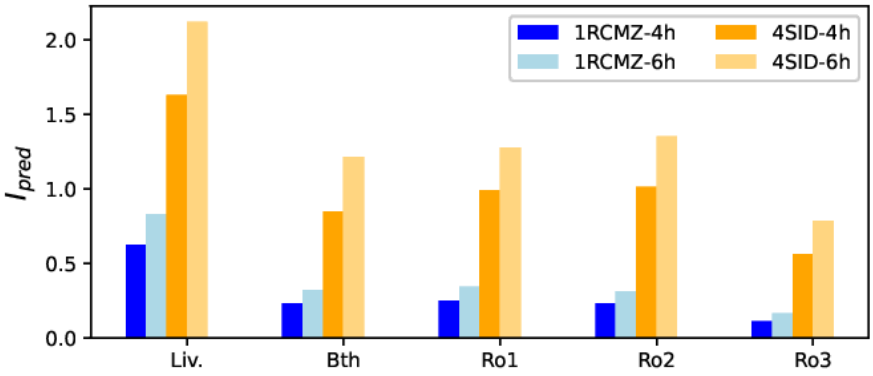

Figure II.13: Prediction quality comparison between different scheme.

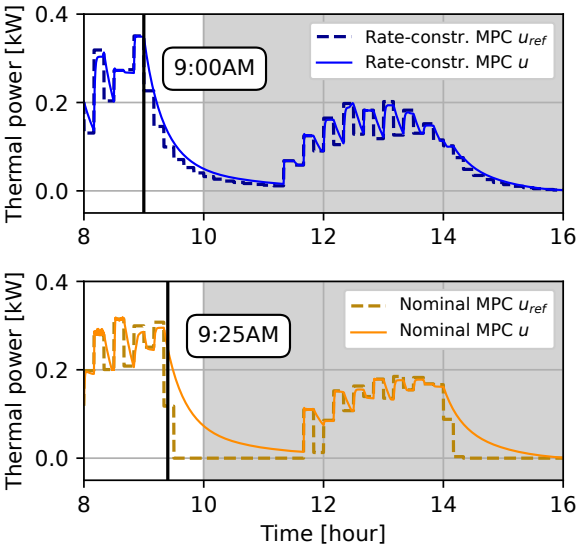

Figure II.14: Radiator power tracking performance. Two subplots compare Rate-constrained MPC-1RCMZ in Optimization Problem II.4 and nominal MPC-1RCMZ in Optimization Problem II.2. The gray zone represents non-occupied time requiring lower temperature demand.

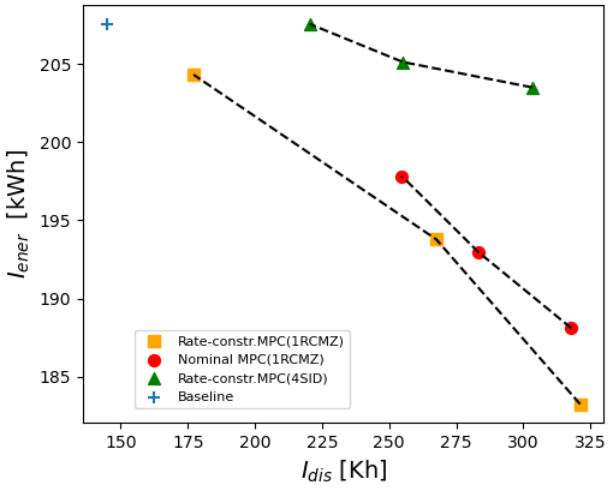

Figure II.15: KPI map for different schemes. I_{dis} is the thermal discomfort indicator, see equation (41). I_{ener} is the energy consumption indicator, see equation (42). For each scheme, three Weighting coefficient w in the MPC is tested. Corresponding value from left to right is $w \in \{100w_0, w_0, 0.1w_0\}$

Figure II.14 shows the effect of the power decrease constraint. Rate-constrained MPC in Optimization Problem II.4 is constrained by the internal loop model. Nominal MPC in Optimization Problem II.2 has no limit on equipment power decrease. Rate-constrained MPC aligns better with the actual power response. At 10:00, the residential building changes to non-occupied mode. The nominal MPC provides a power reference that quickly drops to 0. The heating equipment cannot follow immediately,

MPC has to wait for the radiator to exhaust its power. Conversely, Rate-constrained MPC suggests a first-order-shaped curve with a slow descent, adequately planning the usage of residual heat, thereby starting the power decrease process earlier than the nominal MPC. Equipment thermal inertia is then fully managed.

In run-time environment, two Key Performance Indicators (KPI) are defined to evaluate the control scheme. These KPI rate the performance of energy consumption and discomfort according to [1]. The thermal discomfort indicator is

$$I_{dis} = \sum_{i=1}^{n} \sum_{k=1}^{n_{data}} max(T_{set,i}\left(t_k\right) - T_i\left(t_k\right), 0) \tag{41}$$

where $max()$ corresponds to a saturation on temperature error, only under-heating error is counted as discomfort. $T_{set,i}$ is the heating set-point temperature of each zone. T_i is the measured temperature of each zone. The unit of this indicator is converted to $[Kelvin \cdot hour]$ according to simulation sampling time.

The consumption indicator is

$$I_{ener} = \sum_{i=1}^{n} \sum_{k=1}^{n_{data}} P_i\left(t_k\right) \tag{42}$$

where P_i is heating power of radiator in each zone. The unit of this indicator is converted to $kWatt \cdot hour$ according to simulation sampling time.

Figure II.15 synthesizes and compares the KPI of several schemes. Including original baseline control scheme in the testcase, a set of hierarchical control scheme with Optimization Problem II.2 and Optimization Problem II.4. The better performing scheme will be closer to the bottom left corner. However, compromises between the two objectives will shift the result from the upper left to the lower right, depending on the objective weights. It corresponds to the Parento front of the optimization. The comparison between the nominal MPC and the Rate-constrained MPC is interesting. With a heavy weight on consumption ($w = 100w_0$), Rate-constrained MPC consumes less energy without significantly decreasing thermal comfort. It fully used the residue heat in equipment. When the consumption weight is relatively low ($w = w_0$), the results of the two models are comparable. Adding thermal power rate constraint does not absolutely imply a reduction in energy consumption. Model errors also fluctuate results. The residue heat released by equipment sometimes does not lead to a overheat, but covers the under-heating at beginning of non-occupied time. This happens in low radiation days, such as January 23 midday in Liv., see figure II.11. Thus, compensated model will not result in noticeable improvements when the energy savings target is not important enough. With 4SID model, Rate-constrained MPC has more discomfort, as the 4SID model's divergence leads to considerable overheating. The 1RCMZ model achieves approximately 3–12% reduction in energy consumption, while having comparable or even lower levels of thermal discomfort compared to the 4SID model.

4.4 Summary

Both the 1RCMZ and 4SID models effectively capture the thermal behavior of individual zones within a building. The 1RCMZ model exhibited stable performance in the dataset length sensitivity test. Although the data validation performance is comparable for both models, the RC and 4SID models

demonstrate different closed-loop performances when integrated into the MPC framework. The RC model maintains good closed-loop performance when trained on historical datasets. The proposed methodology, combining MPC and VRFT, enables efficient, and even optimized, management of the produced power for thermal regulation purposes. Furthermore, the integration of power variation constraints improves energy management by enhancing the trade-off between energy consumption and occupant comfort.

5 Chapter Conclusion and Outlook

This chapter investigated hierarchical MPC for building thermal management, comparing gray-box and black-box modeling approaches under limited historical operational data conditions. The research addressed a question: how do modeling choices between physics-informed and purely data-driven approaches affect control performance when training data are limited?

5.1 Conclusion

The gray-box RC model and black-box 4SID model were compared through both open-loop prediction validation and closed-loop control performance evaluation. The study revealed that while both modeling approaches achieve comparable open-loop prediction accuracy, their closed-loop performance diverges significantly. The RC model demonstrated superior extrapolation capabilities by incorporating physical structure, achieving 3-12% energy savings compared to the 4SID model while maintaining comparable or lower thermal discomfort levels. This performance advantage stems from the RC model's ability to enforce energy conservation principles through its network topology, enabling reliable predictions even when the MPC optimization explores operating regions distant from training data.

The proposed hierarchical control framework successfully addressed two key challenges in building thermal management: equipment diversity and limited data availability. At the supervisory level, rate-constrained MPC formulation explicitly accounts for asymmetric power dynamics in high-thermal-inertia HVAC equipment, with constraints identified directly from operational data. The rate constraint enables the MPC to anticipate gradual thermal power evolution, optimizing the utilization of residual heat and improving energy management. At the tracking level, VRFT provides a model-free approach to tune decentralized PID controllers without requiring explicit HVAC system models, effectively handling equipment heterogeneity across different zones.

The validation results on the MZRH testcase demonstrated several important findings. First, the RC model exhibits deterministic behavior in dataset length sensitivity tests, achieving good performance with as little as one day of training data, while the 4SID model requires approximately ten days to reach comparable accuracy. Second, the 4SID model shows greater sensitivity to dataset selection and operating conditions, occasionally producing outlier performance even with apparently sufficient data length. Third, the rate-constrained MPC formulation provides consistent benefits when energy consumption minimization is prioritized, fully exploiting equipment thermal inertia for improved efficiency.

The comparison between gray-box and black-box modeling reveals complementary trade-offs. The RC model leverages physical structure to achieve rapid convergence and consistent performance with minimal data, naturally enforcing energy conservation through its topology. The 4SID model

provides greater flexibility in approximating complex dynamics without requiring physical insight, but demands more extensive training data and exhibits higher sensitivity to dataset characteristics. For systems where physical understanding is available and reliable, gray-box approaches offer practical advantages.

5.2 Outlook: Model-free Predictive Control with DeePC

The findings of this chapter motivate the exploration of fully model-free predictive control approaches. While both RC and 4SID models require explicit system identification procedures—whether constrained by physical structure or purely empirical—DeePC. DeePC offers a different paradigm. DeePC bypasses explicit model identification entirely by directly using input-output data trajectories for prediction, based on Willems' fundamental lemma from behavioral systems theory.

The next chapter extends the hierarchical control framework to integrate DeePC at the supervisory level. Several key questions will be addressed: Can DeePC achieve comparable or superior performance to model-based MPC without any modeling process? How does DeePC's direct data-driven prediction compare to the indirect data-driven approaches examined in this chapter? What are the practical considerations for DeePC deployment, including data requirements, hyperparameter selection, and computational complexity?

The hierarchical framework established in this chapter provides an ideal testbed for DeePC evaluation. By maintaining the same power-based separation and VRFT-tuned tracking controllers, the comparison can isolate the impact of replacing model-based prediction with direct data-driven prediction. The DeePC formulation will be validated on the multi-zone office building testcase with air handling systems, extending the methodology to a different building type and HVAC configuration. This progression from model-based to model-free predictive control represents a natural evolution toward reducing the modeling burden while maintaining or improving control performance for practical building thermal management applications.

Chapter III

Model-free Predictive Control with DeePC

1 Preliminaries on Data-driven Control 79
1.1 Direct Data-driven Control Framework 80
1.2 Data-enabled Predictive Control 83
1.3 Subspace Predictive Control 87
1.4 Comparison of DeePC and SPC 90
2 DeePC Implementation and MOSA Validation 92
2.1 Hierarchical DeePC Implementation Setup 93
2.2 Closed-loop Performance and Comparative Analysis 95
2.3 Sensitivity Analysis of DeePC Hyperparameters 99

1 Preliminaries on Data-driven Control

Data-driven control refers to methodologies where measured data participates in control system deployment. Data-driven approaches are categorized into several classes based on how measured data participates in the control design process. Some methods use data to identify parametric models offline before controller synthesis, while others integrate data directly into online control computations. When combined with predictive control—which optimizes control actions over a future horizon—direct data-driven control yields direct data-driven predictive control. DeePC represents one approach within this intersection.

Data-enabled Predictive Control (DeePC) represents a novel approach built upon well-established theoretical foundations in behavioral systems theory. This section identifies the position of DeePC within the broader control methodology landscape. This investigation leads to a discussion of the relationship between DeePC and another popular method, Subspace Predictive Control (SPC). Figure III.1 illustrates the methodological landscape where DeePC originates. The terminology in this diagram is developed throughout the following subsections.

Indirect data-driven predictive control first identifies parametric models from data, then uses these models for prediction and optimization. SPC exemplifies this indirect approach. While both DeePC and SPC belong to data-driven predictive control, they differ in how they leverage measured data: DeePC combines historical trajectories directly, whereas SPC extracts and applies reduced-order parametric representations. For building thermal management, these differences have practical implications: DeePC's model-free characteristic addresses equipment diversity by eliminating supervisory-level modeling requirements, while SPC's compact parametric form offers computational efficiency suitable for embedded building management systems. This distinction has important implications for controller design, computational structure, and adaptation mechanisms, which this chapter investigates.

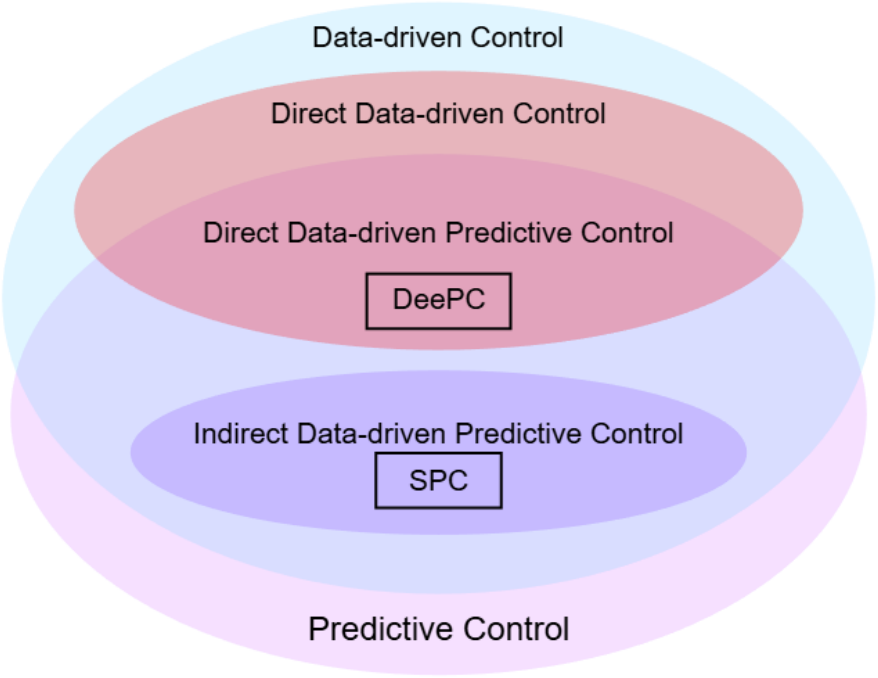

Figure III.1: Venn diagram of data-driven predictive control approaches

1.1 Direct Data-driven Control Framework

The focus of this work, DeePC method, is based on the Willems' fundamental lemma-based direct data-driven control framework. This framework designs controllers by using input-output data directly in the control computation without requiring explicit parametric model identification. For building thermal management, this approach offers a promising solution to equipment diversity by eliminating the need for detailed HVAC system models in the supervisory control layer.

The theoretical foundations of Willems'-lemma-based data-driven control can trace back to behavioral systems theory, introduced by [101]. This framework describes a dynamical system as a set of trajectories that satisfy the system's laws, rather than through parametric representation. Behavioral systems theory has gained renewed interest with data-driven prediction. It is well-suited for this purpose due to its trajectory-based characterization [102].

The key mathematical result is the fundamental lemma [44], which establishes when a finite-length trajectory of a linear time-invariant system fully characterizes the system's finite-horizon behavior. The introduction of the direct data-driven control foundation is based on [103, 104]

Consider a controllable and observable linear time-invariant (LTI) system:

$$\begin{aligned} x[k+1] &= Ax[k] + Bu[k], \\ y[k] &= Cx[k] + Du[k], \end{aligned} \tag{43}$$

where $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{R}^m$ is the input, and $y \in \mathbb{R}^p$ is the output. The matrices A, B, C, D define the system dynamics.

The input-output response over time interval $[0, t-1]$ is:

$$\begin{bmatrix} u_{[0,t-1]} \\ y_{[0,t-1]} \end{bmatrix} = \begin{bmatrix} I_{tm} & \mathbf{0}_{tm\times n} \\ \mathcal{T}_t & \mathcal{O}_t \end{bmatrix} \begin{bmatrix} u_{[0,t-1]} \\ x_0 \end{bmatrix}, \tag{44}$$

where $x_0 \in \mathbb{R}^n$ is the initial state, and the notation $u_{[0,t-1]} := [u[0]^\top, u[1]^\top, \ldots, u[t-1]^\top]^\top$ stacks the input sequence into a vector. The output sequence $y_{[0,t-1]}$ is defined through the Toeplitz matrix $\mathcal{T}_t$ and observability matrix $\mathcal{O}_t$:

$$\mathcal{T}_t := \begin{bmatrix} D & 0 & 0 & \cdots & 0 \\ CB & D & 0 & \cdots & 0 \\ CAB & CB & D & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{t-2}B & CA^{t-3}B & CA^{t-4}B & \cdots & D \end{bmatrix} \in \mathbb{R}^{pt\times mt}, \quad \mathcal{O}_t := \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{t-1} \end{bmatrix} \in \mathbb{R}^{pt\times n}. \tag{45}$$

Equation (44) shows that outputs depend linearly on inputs and the initial state. The Toeplitz matrix $\mathcal{T}_t$ captures how inputs affect outputs. Its lower-triangular structure means outputs depend only on current and past inputs, not future ones. The observability matrix $\mathcal{O}_t$ shows how the initial state influences outputs over time.

We now introduce the Hankel matrix from measured data. The Hankel matrix is denoted by a capital letter with three subscripts: $X_{a,L,M}$ for a time sequence x. The first subscript a indicates the starting time index of the sequence, which appears at the top-left position of the matrix. The second and third subscripts L and M specify the matrix dimensions: height (number of rows) and width (number of columns), respectively. Given a time series $x_d := \{x_d[k]\}_{k=0}^{T_d-1}$, the Hankel matrix with height h is defined as:

$$X_{0,h,T_d-h+1} := \begin{bmatrix} x_d[0] & x_d[1] & \cdots & x_d[T_d-h] \\ x_d[1] & x_d[2] & \cdots & x_d[T_d-h+1] \\ \vdots & \vdots & \ddots & \vdots \\ x_d[h-1] & x_d[h] & \cdots & x_d[T_d-1] \end{bmatrix} \in \mathbb{R}^{h\times(T_d-h+1)}. \tag{46}$$

Each column contains a trajectory segment of length h, forming a sliding window over the time sequence. These columns span the behavioral space of the system, with each column representing a distinct trajectory realization.

Let $u_d := \{u_d[k]\}_{k=0}^{T_d-1}$ and $y_d := \{y_d[k]\}_{k=0}^{T_d-1}$ be input and output sequences of length T_d from an experiment. The vertical concatenated Hankel matrices are constructed as:

$$\begin{bmatrix} U_{0,h,T_d-h+1} \\ Y_{0,h,T_d-h+1} \end{bmatrix} := \left[\begin{array}{cccc} u_d[0] & u_d[1] & \cdots & u_d[T_d-h] \\ u_d[1] & u_d[2] & \cdots & u_d[T_d-h+1] \\ \vdots & \vdots & \ddots & \vdots \\ u_d[h-1] & u_d[h] & \cdots & u_d[T_d-1] \\ \hline y_d[0] & y_d[1] & \cdots & y_d[T_d-h] \\ y_d[1] & y_d[2] & \cdots & y_d[T_d-h+1] \\ \vdots & \vdots & \ddots & \vdots \\ y_d[h-1] & y_d[h] & \cdots & y_d[T_d-1] \end{array}\right]. \tag{47}$$

Each column of the Hankel matrix contains a trajectory segment of length h starting at different times. Similarly to equation (44), the Hankel matrices satisfy:

$$\begin{bmatrix} U_{0,h,T_d-h+1} \\ Y_{0,h,T_d-h+1} \end{bmatrix} = \begin{bmatrix} I_{hm} & \mathbf{0}_{hm\times n} \\ \mathcal{T}_h & \mathcal{O}_h \end{bmatrix} \begin{bmatrix} U_{0,h,T_d-h+1} \\ X_{0,1,T_d} \end{bmatrix}, \tag{48}$$

where

$$X_{0,1,T_d} = \begin{bmatrix} x_d[0] & x_d[1] & \cdots & x_d[T_d-1] \end{bmatrix} \tag{49}$$

contains the state samples. Subscript d emphasizes data from experiments.

A key concept is persistently exciting signals.

Persistency of Excitation [44] *The signal $u_{[0,T_d-1]} \in \mathbb{R}^m$ is persistently exciting of order h if the matrix*

$$U_{0,h,T_d-h+1} = \begin{bmatrix} u[0] & u[1] & \cdots & u[T_d-h] \\ u[1] & u[2] & \cdots & u[T_d-h+1] \\ \vdots & \vdots & \ddots & \vdots \\ u[h-1] & u[h] & \cdots & u[T_d-1] \end{bmatrix} \in \mathbb{R}^{mh\times(T_d-h+1)} \tag{50}$$

has full row rank mh.

Persistency of excitation guarantees that the input signal explores the input space over a sufficient time horizon. This condition is essential for constructing informative Hankel matrices that capture complete system dynamic. For building thermal systems, satisfying this condition requires careful experimental design, as operational data from closed-loop control typically exhibits limited excitation. Dedicated identification experiments with varying thermal power inputs across different zones can provide the necessary signal richness while respecting occupant comfort constraints.

Fundamental Lemma [44]. *Consider a controllable linear time-invariant system of the form (43) with state dimension n. Let $u_d \in \mathbb{R}^{m\times T_d}$ and $y_d \in \mathbb{R}^{p\times T_d}$ be input-output data collected from this system. If the input sequence u_d is persistently exciting of order mh, then any trajectory $[u^\top, y^\top]^\top$ of length $(m+p)h$ belongs to the system behavior if and only if there exists a vector $g \in \mathbb{R}^{T_d-h+1}$ such that*

$$\begin{bmatrix} U_{0,h,T_d-h+1} \\ Y_{0,h,T_d-h+1} \end{bmatrix} g = \begin{bmatrix} u \\ y \end{bmatrix}, \tag{51}$$

where U_{0,h,T_d-h+1} and Y_{0,h,T_d-h+1} are Hankel matrices constructed from the data sequences.

This lemma provides a non-parametric representation through the column space of data-constructed Hankel matrices, showing that persistently exciting data reveals system dynamics without explicit identification of state-space matrices. The persistency of excitation condition ensures the input signal explores all relevant directions in the system dynamics of equation (48) [105].

Beyond trajectory prediction, the fundamental lemma enabled direct controller synthesis without explicit system identification. The data-driven framework retains classical open-loop and closed-loop representations by replacing analytical models with equation (51). This approach preserves the linear mathematical structure, enabling direct application of Lyapunov stability analysis, linear matrix inequalities, and optimal control [103].

Several data-driven control methods emerged from this framework. Data-driven stabilization showed that classical control problems can be reformulated [103], including state feedback, output

feedback, and linear quadratic regulation. These concepts were extended to finite-horizon linear quadratic regulation via semidefinite programming [106]. Robustness to measurement noise was addressed through matrix S-lemma approaches [107]. Frameworks for combining first-principles models with measured data showed the potential of prior models informed direct data-driven predictors [108]. Extensions to nonlinear systems through kernel methods and local linearization have been pursued [109], alongside formulations for linear parameter-varying systems [110]. Behavioral theory for stochastic systems connected the fundamental lemma with polynomial chaos expansions for data-driven stochastic optimal control [111].

This rich methodological gallery provides several advantages for building thermal management. The direct data-driven framework further addresses system diversity by eliminating the need for explicit supervisory control models across different HVAC configurations [99, 58, 112]. The framework provides established plugins for both predictive control and adaptive control, enabling systematic controller design and online adaptation. The potential for adaptive mechanisms is particularly relevant for building systems, where time-varying factors such as seasonal weather patterns, occupancy variations, and equipment aging require continuous controller adjustment. These adaptive capabilities form the focus of chapter IV, where online data matrix updating strategies are investigated to maintain control performance under evolving operating conditions.

1.2 Data-enabled Predictive Control

Direct data-driven predictive control employs equation (51) as the system representation for prediction. Rather than deriving an explicit parametric model, this approach uses the fundamental lemma constraint directly within a receding-horizon optimization framework to predict future system behavior.

The standard formulation of direct data-driven predictive control (DDPC) is:

Optimization Problem III.1 Direct Data-driven Predictive Control

$$\min_{g} \quad J(\mathbf{y}, \mathbf{u}) + \lambda_g ||g||_p \tag{52a}$$

$$\text{s.t.} \quad \begin{bmatrix} U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} \mathbf{u} \\ \mathbf{y} \end{bmatrix}, \tag{52b}$$

$$\mathbf{u} \in \mathcal{U}^N, \quad \mathbf{y} \in \mathcal{Y}^N, \tag{52c}$$

$$\text{with} \quad U_f := U_{0,N,T_d-N+1}, \tag{52d}$$

$$Y_f := Y_{0,N,T_d-N+1} \tag{52e}$$

---

where $J(\mathbf{y}, \mathbf{u})$ is the economic objective like energy consumption minimization and thermal comfort regulation. $\mathbf{u}$ and $\mathbf{y}$ are the predicted input and output trajectories over the prediction horizon of N steps, $\mathcal{U}^N$ and $\mathcal{Y}^N$ denote the constraint sets for inputs and outputs, $g \in \mathbb{R}^{T_d-N+1}$ is the trajectory combination vector that indexes and linearly combines historical trajectories stored in the Hankel matrices U_f and Y_f, and λ_g is the regularization weight on g (typically using ℓ_1 or ℓ_2 norms). The constraint (52b) implements Willems' fundamental lemma as the prediction mechanism. In nominal MPC, an explicit model predicts output trajectories from input decisions; here, the fundamental lemma

constraint serves this role by ensuring predicted trajectories lie in the behavioral space characterized by historical data.

DDPC provides a representative formulation, embedding the fundamental lemma's trajectory representation within a receding-horizon optimal control framework [45]. DDPC formulates predictions and control action computation as a single optimization where decision variables include both future control inputs and coefficients for combining historical trajectories. The optimization objective typically balances tracking performance, control effort, and regularization terms that perform trajectory combinations. Constraints enforce that the predicted trajectory lies in the behavioral space characterized by the data matrix while satisfying operational limits on inputs and outputs. The term "model-free predictive control" sometimes appears in literature, though "modeling-free" more accurately describes the elimination of the parametric identification workflow rather than the absence of prediction model.

Data-enabled Predictive Control (DeePC) extends this formulation by incorporating initial condition estimation through past data windows. The DeePC optimization is:

Optimization Problem III.2 Data-enabled Predictive Control

$$\min_{g,\sigma_y} \quad J(\mathbf{y},\mathbf{u}) + \lambda_y ||\sigma_y||_2^2 + \lambda_g ||g||_p \tag{53a}$$

$$\text{s.t.} \begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{ini} \\ \mathbf{y}_{ini} \\ \mathbf{u} \\ \mathbf{y} \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_y \\ 0 \\ 0 \end{bmatrix}, \tag{53b}$$

$$\mathbf{u} \in \mathcal{U}^N, \quad \mathbf{y} \in \mathcal{Y}^N, \tag{53c}$$

$$\text{with} \quad \begin{bmatrix} U_p \\ U_f \end{bmatrix} := U_{0,T_{ini}+N,T_d-T_{ini}-N+1}, \tag{53d}$$

$$\begin{bmatrix} Y_p \\ Y_f \end{bmatrix} := Y_{0,T_{ini}+N,T_d-T_{ini}-N+1} \tag{53e}$$

---

where U_p, Y_p denote the past Hankel matrices (first T_{ini} rows) for initial condition estimation, U_f, Y_f denote the future Hankel matrices (last N rows) for trajectory prediction, $g \in \mathbb{R}^{T_d-T_{ini}-N+1}$ is the trajectory combination vector, and $\mathbf{u}_{ini}, \mathbf{y}_{ini}$ are the most recent measured input and output sequences of length T_{ini} establishing the current initial condition, while $\mathbf{u}, \mathbf{y}$ are the predicted input and output trajectories over horizon N.

The slack variable σ_y with weight λ_y relaxes the past output matching constraint to ensure feasibility under measurement noise and model mismatch. The regularization term $\lambda_g ||g||_p$ (with $p \in \{1, 2\}$) penalizes excessive reliance on potentially irrelevant or noisy historical trajectories. The training dataset (u_d, y_d) must be persistently exciting of order $T_{ini} + N$ to satisfy the fundamental lemma requirements.

It is important to note that the DeePC optimization problem equation (III.2) is convex and quadratic. The economic objective $J(\mathbf{y}, \mathbf{u})$ contains quadratic terms for both thermal comfort and energy consumption. Non-quadratic elements, such as thermal discomfort penalties and ℓ_1 regularization, are reformulated via auxiliary variables and soft constraints to maintain convexity. This convexity guarantees numerical stability and global optimality at each control step, enabling the use of

efficient quadratic programming solvers such as qpOASES for reliable real-time implementation. The hierarchical control architecture facilitates this convexity by decoupling thermal power computation from equipment regulation, ensuring that the supervisory-level optimization remains tractable for building management systems with limited computational resources.

The trajectory predictor is formulated based on the fundamental lemma:

$$\begin{bmatrix} U_{0,T_{ini},T_d-T_{ini}-N+1} \\ Y_{0,T_{ini},T_d-T_{ini}-N+1} \\ U_{T_{ini},N,T_d-T_{ini}-N+1} \\ Y_{T_{ini},N,T_d-T_{ini}-N+1} \end{bmatrix} g = \begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{ini} \\ \mathbf{y}_{ini} + \sigma_y \\ \mathbf{u} \\ \mathbf{y} \end{bmatrix}, \tag{54}$$

This formulation distinguishes the Hankel matrix into two parts, denoted by subscripts p (past) and f (future) in equation (54). This partition assigns distinct roles to each segment. The past Hankel matrices U_p and Y_p encode historical input-output behavior that establishes the initial condition for prediction. The future Hankel matrices U_f and Y_f capture the system's response over the prediction horizon. Through the trajectory combination vector g, the formulation selects and combines historical trajectory segments to predict future behavior consistent with current initial conditions.

The estimation of current initial condition can be interpreted through the underlying system structure. The Hankel matrices embed the Toeplitz and observability matrices from equation (48). Specifically, for LTI systems, we can decompose:

$$Y_p = \mathcal{T}_{T_{ini}} U_p + \mathcal{O}_{T_{ini}} X_p, \tag{55a}$$
$$Y_f = \mathcal{T}_N U_f + \mathcal{O}_N X_f, \tag{55b}$$

where $X_p \in \mathbb{R}^{n\times(T_d-T_{ini}-N+1)}$ contains the initial states of trajectories in the dataset, and X_f contains states at the beginning of future segments.

The controllability matrix $\mathcal{C}_L$ relates initial states across time horizons:

$$\mathcal{C}_L := \begin{bmatrix} A^{L-1}B & A^{L-2}B & A^{L-3}B & \cdots & B \end{bmatrix} \in \mathbb{R}^{n\times mL}. \tag{56}$$

The state propagation over T_{ini} steps follows:

$$X_f = A^{T_{ini}} X_p + \mathcal{C}_{T_{ini}} U_p. \tag{57}$$

This decomposition reveals the implicit model structure underlying the trajectory representation (54). The vector g performs a weighted combination of historical trajectory columns, where each column encapsulates system behavior under specific initial conditions and input sequences. For LTI systems, the Toeplitz and observability matrices remain constant across all columns, ensuring that any linear combination via g produces a trajectory consistent with the underlying dynamics in equation (43).

Substituting (55a), (55b), and (57) into (54), and following algebraic manipulations similar to those in subspace identification literature, we obtain the underlying prediction model:

$$\begin{aligned} \mathbf{y} = {} & \mathcal{T}_N \mathbf{u} + \mathcal{O}_N A^{T_{ini}} \mathcal{O}_{T_{ini}}^{\dagger} (\mathbf{y}_{ini} + \sigma_y) \\ & + \left(\mathcal{O}_N \mathcal{C}_{T_{ini}} - \mathcal{O}_N A^{T_{ini}} \mathcal{O}_{T_{ini}}^{\dagger} \mathcal{T}_{T_{ini}} \right) \mathbf{u}_{ini}. \end{aligned} \tag{58}$$

This expression shows that the future output $\mathbf{y}$ depends on: (1) the future input $\mathbf{u}$ through the Toeplitz matrix, (2) the past output $\mathbf{y}_{ini}$ through the observability matrices that encode state estimation

and propagation, and (3) the past input $\mathbf{u}_{ini}$ through combined controllability and observability effects. The superscript † denotes the Moore-Penrose pseudoinverse.

With equation (55) and equation (58), we can derive the expression of g in equation (54). Rewrite equation (55) according to the matrix organization in equation (54):

$$\begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} = \Gamma \begin{bmatrix} X_{0,1,T_d} \\ U_p \\ U_f \end{bmatrix}, \tag{59a}$$

$$\text{with} \quad \Gamma = \begin{bmatrix} \mathbf{0} & I_{T_{ini}m} & \mathbf{0} \\ \mathcal{O}_{T_{ini}} & \mathcal{T}_{T_{ini}} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & I_{Nm} \\ \mathcal{O}_N A^{T_{ini}} & \mathbf{0} & \mathcal{T}_N \end{bmatrix}, \tag{59b}$$

where Γ represents the trajectory dynamics matrix that maps the initial state and input sequences to the input-output trajectory.

Multiplying both sides by g,

$$\Gamma \begin{bmatrix} X_{0,1,T_d} g \\ U_p g \\ U_f g \end{bmatrix} = \Gamma \begin{bmatrix} X_{0,1,T_d} g \\ \mathbf{u}_{ini} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{u}_{ini} \\ \mathbf{y}_{ini} + \sigma_y \\ \mathbf{u} \\ \mathbf{y} \end{bmatrix}, \tag{60a}$$

$$\text{with} \quad U_p g = \mathbf{u}_{ini},\ U_f g = \mathbf{u} \quad \text{in (53b).} \tag{60b}$$

The term $X_{0,1,T_d} g$ represents the weighted combination of initial state trajectories from the data, effectively estimating the initial condition for the online trajectory. Thus, g serves as an initial condition composition vector that reconstructs the current system behavior from historical data.

When the LTI system is persistently excited and the Hankel matrix has full row rank, g can be explicitly computed as a least square solution:

$$g = \left(\begin{bmatrix} Y_p \\ U_p \\ U_f \end{bmatrix}^\top \begin{bmatrix} Y_p \\ U_p \\ U_f \end{bmatrix} \right)^{-1} \begin{bmatrix} Y_p \\ U_p \\ U_f \end{bmatrix}^\top \begin{bmatrix} \mathbf{y}_{ini} + \sigma_y \\ \mathbf{u}_{ini} \\ \mathbf{u} \end{bmatrix}. \tag{61}$$

The underlying model in equation (58) is optimized with economic objectives in DeePC. The slack variable σ_y allows this model to deviate from current measurements when such deviation reduces overall cost through improved economic performance. This coupling introduces a distinctive feature: the optimization prioritizes trajectory regions aligned with economic objectives. By jointly optimizing g and σ_y, DeePC dynamically emphasizes behavioral space regions that offer economic benefits.

This feature presents both opportunities and challenges for building thermal management. On one hand, DeePC naturally focuses on economically relevant operating regimes. On the other hand, the economic optimum shifts continuously with external conditions. For building systems, solar radiation variations alter heat gain profiles throughout the day, affecting the optimal thermal power distribution across zones. Weather-driven changes in envelope heat transfer coefficients modify the building's thermal response characteristics. Occupancy patterns introduce step changes in internal gains and comfort requirements. These factors disturb the implicit model representing thermal power response

to control actions. Furthermore, this dynamic modeling characteristic complicates regularization weight tuning. The parameters λ_y and λ_g must balance competing objectives: economic performance, prediction accuracy, and robustness to external variations. These weights determine how aggressively the controller adapts its implicit model to pursue economic gains versus maintaining prediction consistency across varying operating conditions.

Several theoretical extensions of DeePC have addressed robustness and operational challenges. Regularized DeePC incorporated norm-based penalties on trajectory combination coefficients, with theoretical justification through distributionally robust stochastic optimization [113]. Chance-constrained formulations explicitly handled output probability requirements [114]. Extensions to networked systems enabled decentralized implementations [115, 116], while adaptive variants incorporated online data matrix updates [102].

1.3 Subspace Predictive Control

The Indirect data-driven predictive control uses data as an intermediate step to obtain a parametric model. While avoiding detailed first-principles derivations, this approach constructs explicit state-space representations from measured trajectories before performing predictions. Subspace predictive control exemplifies this methodology, combining subspace identification with MPC frameworks [117, 118].

Subspace identification theory established foundations for estimating state-space models directly from input-output data without requiring knowledge of system structure [93, 16]. Integration with predictive control led to Subspace Predictive Control (SPC) algorithms that simultaneously identify and control unknown systems [117, 118].

SPC is an indirect data-driven approach: first identifying a parametric predictor offline from Hankel-structured data, then using this fixed model for online optimization. The SPC optimization formulation [117] is:

Optimization Problem III.3 Subspace Predictive Control with One-step-ahead Prediction

$$\min_{\mathbf{u},\mathbf{x}} \quad J(\mathbf{y},\mathbf{u}) \tag{62a}$$

$$\text{s.t.} \quad \forall j = 0,\dots,N-1, \tag{62b}$$

$$x[k+j+1] = \hat{A}x[k+j] + \hat{B}u[k+j], \tag{62c}$$

$$y[k+j+1] = \hat{C}x[k+j+1] + \hat{D}u[k+j], \tag{62d}$$

$$x_0 = x[k] \tag{62e}$$

$$\mathbf{u} \in \mathcal{U}^N, \quad \mathbf{y} \in \mathcal{Y}^N, \tag{62f}$$

where $J(\mathbf{y},\mathbf{u})$ is the economic objective, $\mathbf{u} = [u[k],\dots,u[k+N-1]]^\top$ and $\mathbf{y} = [y[k],\dots,y[k+N-1]]^\top$ are the predicted input and output trajectories over horizon N, and $\mathbf{x} = [x[k],\dots,x[k+N]]^\top$ is the predicted state trajectory. The state-space matrices $(\hat{A},\hat{B},\hat{C},\hat{D})$ are identified offline through subspace identification methods. The model provides one-step-ahead predictions, with the index $[k+j]$ denoting the time step within the prediction horizon.

The multiple-shooting approach is commonly used in implementation for computational efficiency, but the formulation in Optimization Problem III.3 presents the dynamics in sequential form for clarity and simplicity. Both formulations are equivalent in optimality but differ in computational structure.

To connect one-step-ahead prediction with trajectory prediction, consider the state-space model in equation (43). The multi-step prediction over horizon N can be written in stacked form:

$$\begin{bmatrix} y[k] \\ y[k+1] \\ \vdots \\ y[k+N-1] \end{bmatrix} = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{N-1} \end{bmatrix} x[k] + \begin{bmatrix} D & 0 & \cdots & 0 \\ CB & D & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ CA^{N-2}B & CA^{N-3}B & \cdots & D \end{bmatrix} \begin{bmatrix} u[k] \\ u[k+1] \\ \vdots \\ u[k+N-1] \end{bmatrix}, \quad (63)$$

or compactly defined as equation (45). This trajectory prediction form is more directly comparable to the DeePC optimization problem and reveals the structural similarity between the two approaches.

The initial state $x[k]$ in equation (62c) is determined from measurements at the beginning of each prediction horizon. Substituting the state prediction structure into equation (III.3) yields the trajectory prediction formulation:

Optimization Problem III.4 Subspace Predictive Control with Trajectory Prediction

$$\begin{aligned} \min_{\mathbf{u}} \quad & J(\mathbf{y}, \mathbf{u}) && \text{(64a)} \\ \text{s.t.} \quad & \mathbf{y} = \hat{\mathcal{O}}_N x[k] + \hat{\mathcal{T}}_N \mathbf{u}, && \text{(64b)} \\ & x_0 = x[k] && \text{(64c)} \\ & \mathbf{u} \in \mathcal{U}^N, \quad \mathbf{y} \in \mathcal{Y}^N, && \text{(64d)} \end{aligned}$$

---

where $\mathbf{u} = [u[k], \ldots, u[k+N-1]]^\top$ and $\mathbf{y} = [y[k], \ldots, y[k+N-1]]^\top$ are the predicted input and output trajectories over horizon N, and the predictor matrices $(\hat{\mathcal{O}}_N, \hat{\mathcal{T}}_N)$ are identified offline through subspace identification.

The initial state $x[k]$ can be obtained through various state estimation approaches, such as Kalman filtering for stochastic systems or deterministic observers for noise-free cases. To be comparable with DeePC approach we use the same estimator,

$$x[k] = \Phi_{ini} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \end{bmatrix}, \quad (65)$$

where $\mathbf{y}_{ini} = [y[k-T_{ini}], \ldots, y[k-1]]^\top$ and $\mathbf{u}_{ini} = [u[k-T_{ini}], \ldots, u[k-1]]^\top$ are past measurements of length T_{ini}, and Φ_{ini} is the initial condition estimator matrix identified offline from the same Hankel-structured data used for system identification. Substituting equation (65) into constraint (64b) gives:

$$\mathbf{y} = \hat{\mathcal{O}}_N \Phi_{ini} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \end{bmatrix} + \hat{\mathcal{T}}_N \mathbf{u}. \quad (66)$$

Unlike DeePC's constraint (53b) which uses the full Hankel structure online with variable g, SPC's constraint (64b) employs a compact parametric model extracted from the same data, eliminating the need for online data matrix operations. The initial condition estimation through Φ_{ini} in equation (65) provides a direct data-driven alternative to traditional state observers.

Using the same Hankel partition as DeePC equation (53e), where U_p, Y_p encode past behavior and U_f, Y_f encode future behavior, the future output Hankel satisfies equation (55b).

To estimate $\mathcal{O}_N$ without knowing X_f, subspace identification exploits the orthogonal projection of Y_f onto the orthogonal complement of U_f's row space:

$$\bar{Y}_f := Y_f \Pi_{U_f^\perp} = \mathcal{O}_N X_f \Pi_{U_f^\perp}, \tag{67}$$

where $\Pi_{U_f^\perp} = I - U_f^\dagger U_f$ eliminates the direct feedthrough term $\mathcal{T}_N U_f$. Applying singular value decomposition to $\bar{Y}_f$:

$$\bar{Y}_f = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} \begin{bmatrix} V_1^\top \\ V_2^\top \end{bmatrix}, \tag{68}$$

the system order n is determined by the number of significant singular values in Σ_1. The extended observability matrix and state sequence are then estimated as:

$$\hat{\mathcal{O}}_N = U_1 \Sigma_1^{1/2}, \tag{69a}$$

$$\hat{X}_f = \Sigma_1^{1/2} V_1^\top. \tag{69b}$$

This factorization reveals the role of Hankel matrices: their column space characterizes all trajectories compatible with the underlying LTI system, while the singular value decomposition extracts a dominant n-dimensional subspace that captures the system dynamics.

To construct a predictor that maps past measurements $(\mathbf{y}_{ini}, \mathbf{u}_{ini})$ to the current state, we examine how the past Hankel matrices relate to the latent state sequence. From equation (48), the past data admits the decomposition

$$\begin{bmatrix} Y_p \\ U_p \end{bmatrix} = \begin{bmatrix} \mathcal{O}_{T_{ini}} X_p + \mathcal{T}_{T_{ini}} U_p \\ U_p \end{bmatrix}, \tag{70}$$

which shows that the column space of $\begin{bmatrix} Y_p \\ U_p \end{bmatrix}$ is parameterized by the pair (X_p, U_p).

The future state sequence satisfies the state propagation relation

$$X_f = A^{T_{ini}} X_p + \mathcal{C}_{T_{ini}} U_p, \tag{71}$$

indicating that X_f is also a linear function of (X_p, U_p). Therefore, both X_f and $\begin{bmatrix} Y_p \\ U_p \end{bmatrix}$ depend linearly on the same latent variables, which implies the existence of a linear mapping

$$X_f = \Phi_{ini} \begin{bmatrix} Y_p \\ U_p \end{bmatrix}, \tag{72}$$

provided that the input data is persistently exciting.

Rather than explicitly computing this mapping from system matrices, subspace identification estimates it directly from data via least-squares regression. Using the estimated state sequence $\hat{X}_f$, we define the data-driven state estimator

$$\hat{x}[k] = \Phi_{ini} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \end{bmatrix}, \quad \Phi_{ini} = \hat{X}_f \begin{bmatrix} Y_p \\ U_p \end{bmatrix}^\dagger, \tag{73}$$

which is consistent with the linear relation in (72).

Next, the Toeplitz matrix is estimated from the residual dynamics. From equation (55b), we have

$$Y_f = \hat{\mathcal{O}}_N \hat{X}_f + \mathcal{T}_N U_f, \tag{74}$$

which leads to the least-squares estimate

$$\hat{\mathcal{T}}_N = \left(Y_f - \hat{\mathcal{O}}_N \hat{X}_f\right) U_f^{\dagger}. \tag{75}$$

Substituting the state estimator into the lifted prediction model yields

$$\begin{aligned} \mathbf{y} &= \hat{\mathcal{O}}_N \hat{x}[k] + \hat{\mathcal{T}}_N \mathbf{u} \\ &= \hat{\mathcal{O}}_N \Phi_{ini} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \end{bmatrix} + \hat{\mathcal{T}}_N \mathbf{u}. \end{aligned} \tag{76}$$

Substituting (73) and (75) gives

$$\mathbf{y} = \hat{\mathcal{O}}_N \hat{X}_f \begin{bmatrix} Y_p \\ U_p \end{bmatrix}^{\dagger} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \end{bmatrix} + \left(Y_f - \hat{\mathcal{O}}_N \hat{X}_f\right) U_f^{\dagger} \mathbf{u}. \tag{77}$$

Rearranging terms, we obtain

$$\mathbf{y} = Y_f U_f^{\dagger} \mathbf{u} + \hat{\mathcal{O}}_N \hat{X}_f \left(\begin{bmatrix} Y_p \\ U_p \end{bmatrix}^{\dagger} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \end{bmatrix} - U_f^{\dagger} \mathbf{u} \right). \tag{78}$$

Finally, by exploiting the consistency of Hankel matrices constructed from the same dataset, trajectories admit a unified representation

$$\begin{bmatrix} Y_p \\ U_p \\ U_f \end{bmatrix} g = \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \\ \mathbf{u} \end{bmatrix}, \tag{79}$$

which implies

$$\mathbf{y} = Y_f g = Y_f \begin{bmatrix} Y_p \\ U_p \\ U_f \end{bmatrix}^{\dagger} \begin{bmatrix} \mathbf{y}_{ini} \\ \mathbf{u}_{ini} \\ \mathbf{u} \end{bmatrix}. \tag{80}$$

This establishes that the SPC predictor is equivalent to the Hankel-based trajectory representation used in DeePC. In particular, DeePC implicitly performs state estimation and system realization without explicitly identifying the system matrices (A, B, C, D).

1.4 Comparison of DeePC and SPC

The relationship between direct and indirect data-driven predictive control has been investigated [119, 120], revealing that DeePC and SPC share theoretical connections while differing in implementation and flexibility. The following paragraphs examine these similarities and differences across three key aspects.

Use of Hankel matrices

The connection between DeePC and SPC lies in their shared use of Hankel matrix column spaces. Both methods rely on persistently exciting input-output trajectories contain sufficient information to characterize linear system dynamics. Subspace methods construct parametric models from Hankel data through offline dimension reduction via SVD, extracting a compact n-th order state-space representation. This offline identification enables computational efficiency during online control, as the optimization operates on the reduced-order model rather than the full Hankel structure.

DeePC, in contrast, retains the full Hankel structure and performs modeling alongside control optimization. The trajectory combination vector g is computed at each control step, dynamically emphasizing behavioral space regions aligned with current control objectives. As discussed following equation (58), this online modeling prioritizes trajectory regions that offer economic benefits, adapting the implicit model to pursue cost reduction. While this dynamic modeling could adapt to varying conditions, it increases computational burden compared to SPC's fixed parametric predictor.

Modeling objective

The modeling objectives differ between the two approaches. SPC's subspace identification minimizes the ℓ_2-norm of prediction error through least-squares operations inherent in the pseudoinverse computations in equations (73) and (75).

This quadratic criterion provides computational tractability and analytical solutions but offers limited flexibility in shaping the model's emphasis across different operating regions. Specifically, the offline identification process produces a single fixed model that treats all data uniformly, without the ability to emphasize particular operating conditions or adjust the model's attention based on control objectives.

DeePC's formulation provides greater flexibility through online trajectory selection, allowing the implicit model to adapt based on both regularization structure and economic objectives. The regularization term $\lambda_g||g||_p$ enables different trajectory selection strategies: ℓ_1 regularization promotes sparse combinations, emphasizing few dominant historical trajectories, while ℓ_2 regularization encourages smooth trajectory blending across many data segments. Furthermore, DeePC's coupling of modeling with economic objectives in equation (III.2) enables the implicit model to adapt toward economically relevant operating regions through the optimization variable g, a feature absent in SPC's offline identification where the model structure is fixed prior to control deployment.

Under ideal conditions (deterministic LTI systems, noiseless data, exact persistency of excitation), unregularized DeePC with ℓ_2 norm produces predictions equivalent to SPC [120]. However, this equivalence breaks down under realistic conditions involving noise, regularization, or time-varying dynamics, where the different modeling objectives and their associated flexibility lead to distinct controller behaviors.

Choice of system order

The system order encoded in the models relates directly to Hankel matrix dimensions for both methods. For SPC, the identified state dimension n is determined from the singular value spectrum equation (68), with significant drops indicating the effective system order. The observability matrix $\mathcal{O}_N$ achieves rank n when the future Hankel height satisfies $N \geq n$, while the past Hankel height $T_{ini} \geq n$ ensures sufficient information for state estimation.

For DeePC, although no explicit state-space model is extracted, the underlying system dynamics are implicitly encoded through the column space of Hankel matrices. The fundamental lemma requires persistency of excitation of order $T_{ini} + N$ to fully characterize system behavior. When $T_{ini} \geq n$, the Hankel matrices contain sufficient information to estimate an n-th order system's initial condition for prediction. Larger Hankel dimensions improve robustness but increase computational cost and data requirements ($T_d \geq (m + p) \times (T_{ini} + N)$).

This shared dependence on Hankel configuration reveals the connection: both methods extract system information from the same data structure, with SPC performing explicit dimension reduction to order n, while DeePC implicitly works with the full behavioral space encoded in the data.

The relationship between direct and indirect data-driven predictive control has been investigated [120]. This analysis showed that DeePC and SPC exploit the same column space of Hankel matrices. Under assumptions of deterministic linear time-invariant systems, noiseless data, and exact persistency of excitation, unregularized DeePC and SPC produce equivalent predictions. The study further revealed that regularization choices correspond across methods: quadratic regularization in DeePC relates to regularized least-squares identification in SPC.

Implications for building thermal management

The theoretical comparison reveals practical trade-offs for building control deployment. SPC's offline identification and compact parametric predictor offer computational advantages for resource-constrained building management systems, with optimization complexity scaling with state dimension n rather than dataset length T_d. This efficiency enables faster control cycles and reduces hardware requirements for embedded controllers.

DeePC's online trajectory combination adapt implicit modeling toward current economic objectives, potentially improving energy-comfort trade-offs under varying conditions. The hierarchical framework proposed in this work leverages DeePC's model-free characteristic at the supervisory level to address equipment diversity across different HVAC configurations. However, this flexibility comes at computational cost: DeePC's optimization involves the full Hankel structure with decision variable dimension $T_d - T_{ini} - N + 1$, requiring careful regularization tuning and potentially larger hardware memory.

For building thermal management, these characteristics suggest complementary application scenarios: SPC suits applications prioritizing computational efficiency and stable operating conditions. By integrating modeling and control objectives, DeePC achieves precise predictions near the operating point.

The subsequent sections validate both approaches on building thermal benchmarks to quantify these trade-offs under realistic operating conditions. The following sections apply these methods to building thermal control within the hierarchical framework proposed in chapter I, where data-driven predictive control manages supervisory-level thermal planning while VRFT-tuned PID controllers handle equipment-level tracking.

2 DeePC Implementation and MOSA Validation

While building thermal dynamics are inherently nonlinear, they can be locally represented through LPV frameworks where the scheduling variable captures time-varying operating conditions. The BOPTEST benchmark provides a realistic testbed where such quasi-LPV behavior manifests through

varying weather conditions and occupancy patterns. This motivates our comparative study of SPC and DeePC on this platform, as it allows evaluation of how each method handles the parameter-varying nature embedded within the nonlinear dynamics.

This section evaluates the proposed hierarchical control scheme through experimentation on the BOPTEST framework. The analysis proceeds through four subsections: experimental configuration and problem formulation, closed-loop validation of tracking performance, sensitivity analysis of key hyper-parameters, and comparative evaluation against alternative predictive control paradigms. These studies demonstrate the performance of direct data-driven predictive control for building thermal management while providing practical insights for deployment.

2.1 Hierarchical DeePC Implementation Setup

2.1.1 Control Law Formulation

The DeePC algorithm is based on the Optimization Problem III.2. The regularization of g reduces overfitting and provides robustness. ℓ_1 and ℓ_2 norms are compared in this work. The ℓ_2 norm is implemented in a standard quadratic form. To preserve the quadratic structure of the optimization, the implementation of ℓ_1 norm is approximated by incorporating additional soft constraints

$$\min_{\delta_1} \sum \delta_1 \tag{81a}$$

$$s.t. \quad \delta_1 \geq g, \quad \delta_1 \geq -g. \tag{81b}$$

DeePC replaces the prediction model in conventional MPC by a non-parametric system representation in (53b).

In MOSA case study, input and output vector are

$$u = [P_1, ..., P_5, T_e, G], \tag{82}$$

$$y = [T_1, ..., T_5] \tag{83}$$

where 5 zones' HVAC thermal power P_i, outdoor temperature T_e and solar radiation G are considered as system input signal u. Output signal vector y is constructed by temperature of each zone T_i.

Economic objective J is defined as

$$J = J_{energy} + \lambda_{economic} J_{discomfort} \tag{84a}$$

$$\text{with} \quad J_{discomfort} = ||\delta_2||_2^2, \tag{84b}$$

$$\delta_2 \geq y - \bar{y}, \delta_2 \geq \underline{y} - y, \delta_2 \geq 0, \tag{84c}$$

$$J_{energy} = ||u_P - \underline{u}_P||_2^2 \tag{84d}$$

where $J_{energy}, J_{discomfort}$ are energy consumption and thermal discomfort objective respectfully. Thermal discomfort is defined by a soft constraint in (84b) and (84c). δ_2 is a auxiliary variable. $\bar{y}, \underline{y}$ are comfort temperature limits. $u_P = [P_1, ..., P_5]^\top$ is HVAC thermal power inputs. $\underline{u}_P$ is lower power limit vector. $\lambda_{economic}$ is weighting between two objectives. During simulation, J is optimized over the prediction horizon H. After simulation, J evaluates the entire simulation period as a indicator denoted by I_J.

At the tracking level, decentralized PI controllers regulate HVAC actuators to achieve the power references $P_{i,ref}[t_k]$ provided by the supervisory DeePC. For VAV terminal units controlling hot water valve openings, the PI control law is:

$$v_i[t_k] = K_{P,i}(P_{i,ref}[t_k] - P_i[t_k]) + K_{I,i}\int_0^{t_k}(P_{i,ref}[\tau] - P_i[\tau])d\tau, \tag{85}$$

where v_i denotes the valve opening command for zone i, $K_{P,i}$ and $K_{I,i}$ represent the proportional and integral gains, and $P_i[t_k]$ indicates the measured thermal power output at the current time instant. The PI parameters are synthesized using the VRFT methodology (see Optimization Problem II.5), which identifies optimal gains directly from closed-loop data without requiring explicit HVAC system models. This model-free tuning approach aligns with the hierarchical framework's philosophy of eliminating parametric modeling.

2.1.2 Default Control Solution and System Parameters

This section evaluates the proposed hierarchical control scheme and performs sensitivity analysis on its key configuration and hyperparameters.

Parameters	Values	Description
Simulation		
Test case	MOSA	
dt	30s	Simulation sampling time
Climate	Chicago	Weather profile
Period	Jan.9-Jan.15	Days for validation
Controller		
τ	100s	VRFT Ref. model time cst.
T_{PI}	30s	PI sampling time
K_{aw}	$0.3K_i$	Anti-windup gain
T_{DeePC}	30 minutes	DeePC sampling time
T_{ini}	5	Initial condition est. order
λ_y	1	Weighting on regression term
λ_g	100	Weighting on g regularization
Dataset	Jan.2-Jan.8	History data in Hankel
H	6 hours	DeePC prediction horizon
$\lambda_{economic}$	0.05	Economic obj. weight
Solver	'`qpoases`'	Solver in DeePC

Table III.1: Fixed experiment configuration in case study

The configuration of the experiment is shown in table III.1. The choices of hyper-parameters are documented along with the corresponding results. Considering the differences in thermal response models between buildings and significant cost associated with testing individual ones, DeePC in this study uses historical building operation data to construct Hankel matrix. HVAC system power response experiment is relatively simpler than thermal response experiment in practice, the data used

in VRFT process are obtained via pseudo random binary sequence signal experiment, as referenced in [30]. Ventilation and AHU operation in this experiment remain same with validation.

2.2 Closed-loop Performance and Comparative Analysis

2.2.1 Hierarchical DeePC in closed-loop

The tracking layer controller parameters are shown in Tab. III.2. This result confirms that VRFT can capture the implicit operating conditions within the data. The MOSA system uses Variable Air Volume (VAV) terminal units for zone-level heating control. Each VAV unit receives conditioned air from the central Air Handling Unit and is equipped with a hot water heating coil. The hot water valve opening controls the heating power delivered to each zone. The HVAC system allocates different constant airflows $q_{air,VAV,i}$ to each zone VAV terminal. Higher airflow results in a higher rate of thermal exchange in the VAV unit. Consequently, this leads to higher heating power output for the same hot water valve opening.

Zone	$q_{air,VAV}$	K_P	K_I
1	$1.639\ \mathrm{m}^3/\mathrm{s}$	0.00234	$2.525\mathrm{e}^{-4}$
2	$0.328\ \mathrm{m}^3/\mathrm{s}$	0.01139	$12.481\mathrm{e}^{-4}$
3	$0.345\ \mathrm{m}^3/\mathrm{s}$	0.01116	$11.747\mathrm{e}^{-4}$
4	$0.345\ \mathrm{m}^3/\mathrm{s}$	0.01117	$12.182\mathrm{e}^{-4}$
5	$0.255\ \mathrm{m}^3/\mathrm{s}$	0.01543	$19.551\mathrm{e}^{-4}$

Table III.2: PI parameters for each zone tuned by VRFT

The VRFT training dataset contains valve opening commands and the corresponding thermal power measurements from the VAV units. Airflow $q_{air,VAV,i}$ is not included in this dataset, as it is determined by the supply system and remains constant during operation. However, examining the identified controller gains with respect to airflow provides an intuitive validation of the VRFT results.

VRFT assigns smaller controller gains to zones with larger airflow. This result indicates that VRFT successfully adapts to the operating conditions embedded in the data. Specifically, Zone 1 has the largest airflow and receives the smallest gains, while Zone 5 has the smallest airflow and receives the largest gains. This inverse relationship demonstrates the effectiveness of VRFT in tuning controllers for different thermal dynamics.

Comparing these results with the MZRH case (see table II.1) reveals an important difference in the integral gain K_I. The integral gains in the MOSA system (order of 10^{-4}) are significantly larger than those in the MZRH system (order of 10^{-13} to 10^{-15}). This difference reflects the distinct power response characteristics of the two HVAC systems.

In the MZRH system, radiators exhibit large thermal inertia. Power decrease relies on passive convective cooling to the surrounding air, which is a slow process. During power reduction, the valve saturates at zero for extended periods. Larger integral gains would accumulate significant error during this saturation, creating difficulties for the anti-windup mechanism. Therefore, VRFT assigns very small integral gains to the radiator system.

In contrast, the MOSA system uses air-to-water heat exchangers in VAV units. The direct enthalpy injection through hot airflow enables faster heat transfer in both directions. This faster power response

allows the use of larger integral gains to eliminate steady-state errors within the same time frame. The ratio between K_I values in the two systems (approximately 10^9 to 10^{11}) aligns with the difference in their power response time constants.

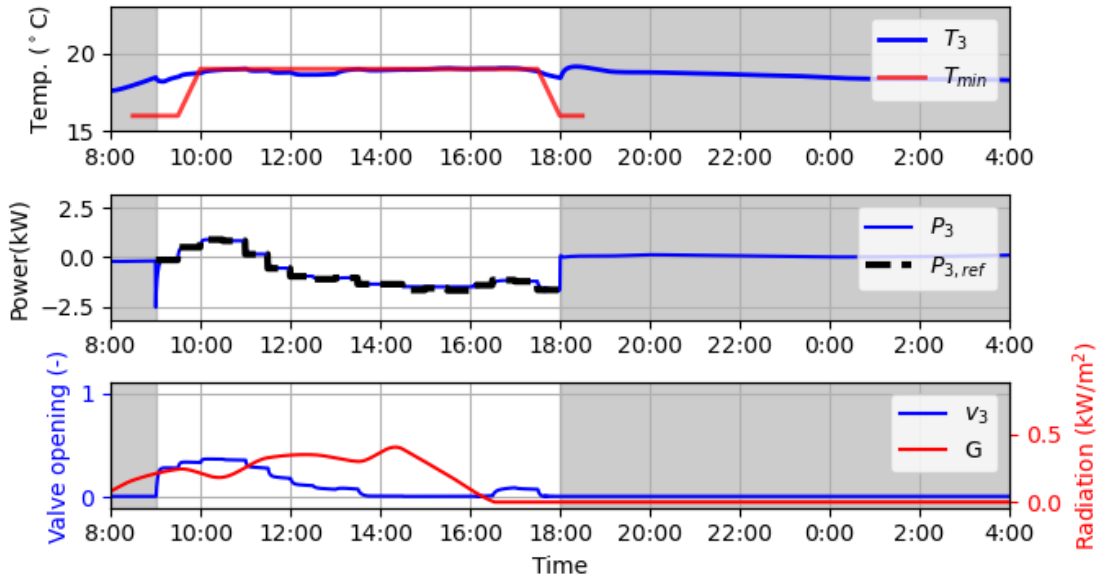

Figure III.2: Simulation result that validates the proposed control scheme. The north zone is presented as a representative example, with similar performance observed across all zones. The weightings are set as $\lambda_y = 1, \lambda_g = 100$. Non-occupied period is marked in gray.

The closed-loop validation of both control layers is illustrated in figure III.2. During the occupied period, the power tracking mechanism demonstrates rapid and accurate tracking of the power reference provided by DeePC. In terms of temperature regulation, the zone temperature is mainly maintained above the discomfort lower limit T_{min}, a boundary defined in $\underline{y}$ in soft constrain (84c).

Since internal heat gain remains relatively constant during the occupied period in this benchmark, solar radiation becomes the most significant disturbance factor. After 11:00, DeePC reduces the VAV power supply and utilizes solar radiation to maintain the zone temperature at the lower bound. This strategy reduces energy consumption while ensuring thermal comfort. At 12:00, DeePC slightly overestimates the heating effect of solar radiation, leading to an underheating of 0.2 °C. However, this deviation is minor and within acceptable limits.

Around 14:00, the zone experiences overheating due to strong solar radiation. At this moment, the VAV valve is completely closed, indicating that the proposed control scheme provides the optimal solution under the given constraints. The negative power values observed after 12:00 reflect the cooling effect from continuous ventilation with cooler outdoor air. This ventilation prevents overheating while maintaining indoor air quality, demonstrating the multi-objective nature of the control strategy.

2.2.2 Parento front comparison

To evaluate the performance of the proposed hierarchical DeePC, a comparison with a conventional hierarchical MPC scheme is presented below. This MPC scheme uses a same modeling methodology

as the previous section, which is a 1st order multi-zone RC model obtained through the Optimization Problem II.1.

The Pareto fronts in figures below demonstrate the performance trade-offs between energy consumption and thermal comfort under varying $\lambda_{economic}$. Better control performance is characterized by lower energy consumption while maintaining low thermal discomfort.

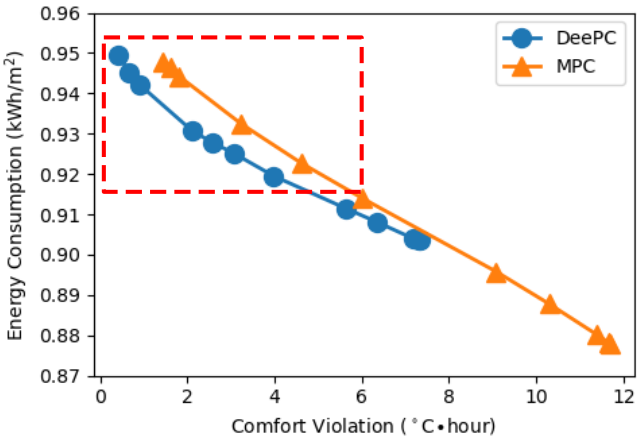

Figure III.3: Pareto fronts comparing DeePC and RC-MPC control schemes under different economic weights.

DeePC vs. RC-MPC Figure III.3 Under identical energy consumption levels in the red dashed line region, DeePC with historical data consistently achieves approximately 1.2 $°\mathrm{C} \cdot \mathrm{hour}$ lower discomfort compared to RC-MPC. When the weight on thermal comfort is significantly high, DeePC reaches lower discomfort levels that are beyond the capability of RC-MPC. This superior performance is attributed to DeePC's data-driven nature, which captures complex thermal dynamics without requiring explicit model structure assumptions. In contrast, the RC model relies on simplified physical representations that may not fully capture all thermal interactions.

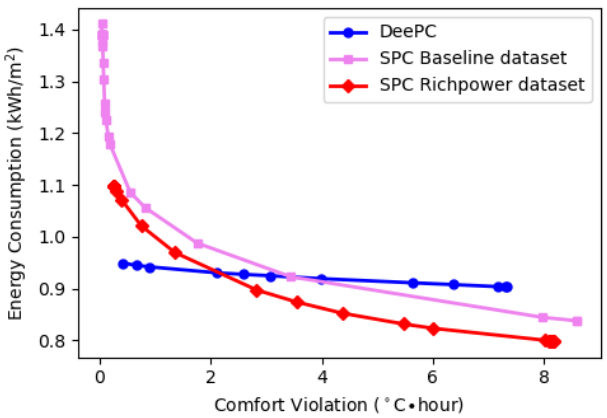

Figure III.4: Pareto fronts comparing DeePC and SPC control schemes with different dataset. Both method initial condition estimation order is 5.

DeePC vs. SPC Figure III.4 The comparison between DeePC and SPC (formulation see the Optimization Problem III.4) reveals differences in their control capabilities. SPC cannot reach some zone at bottom-left, where the comfort violation indicators are very low. When the comfort indicator is low (below 2 $^{\circ}\mathrm{C} \cdot \mathrm{hour}$), SPC's energy consumption rises significantly, indicating overheating issues where the consumption is not well optimized.

This performance difference stems from SPC formulation discussed in section 3. Using the baseline dataset, which consists of normal operation data from Jan.2-Jan.8, both controllers face the challenge of limited excitation. As analyzed, SPC enforces strict equality constraints in its prediction mechanism by assuming the reduced data model fully captures the system behavior. This rigid mechanism requires the input Hankel matrix constructed from U_p to maintain full column rank for accurate predictions. However, when the baseline dataset lacks sufficient excitation—common in normal operations where inputs vary minimally to maintain comfort—the Hankel matrix becomes rank-deficient or ill-conditioned. Consequently, the equality constraints cannot be satisfied accurately, leading to degraded prediction accuracy. This prediction error propagates through the optimization, resulting in sub-optimal control inputs. The observed performance degradation manifests as SPC either overheating the zones (high energy consumption) or failing to maintain comfort (high violation), as evidenced by the steep Pareto front slope in figure III.4.

In contrast, DeePC incorporates regularization terms in its optimization formulation (53b), specifically through slack variables σ_y that relax the equality constraints to $y_f = H_p g + \sigma_y$ and $u_f = H_u g + \sigma_u$. As discussed in Section 1, these slack variables serve two critical functions. First, they accommodate prediction errors arising from rank deficiency by allowing deviations from the strict data-driven prediction. Second, the penalty weights λ_y and λ_u provide a trade-off between fitting the historical data pattern and achieving the control objectives. When the baseline data is insufficient, DeePC can still generate feasible predictions by relying more on the economic and comfort objectives rather than strictly enforcing data consistency. This flexibility is confirmed by the experimental results: DeePC maintains a smooth, consistent Pareto front across all comfort levels, demonstrating stable performance even when the training data does not satisfy persistency of excitation conditions. The approximately 1 $^{\circ}\mathrm{C} \cdot \mathrm{hour}$ comfort improvement at equivalent energy consumption directly validates the theoretical advantage of regularization in handling data-limited scenarios.

The inclusion of SPC with the "Richpower" dataset further validates this analysis. The "Richpower" dataset, generated through pseudo-random binary sequence excitation, provides better input excitation and improves the rank condition of the Hankel matrix. As expected, SPC performance improves significantly, achieving lower energy consumption across most comfort levels. However, SPC with "Richpower" data still cannot match DeePC performance. This remaining gap highlights a key advantage of DeePC: it optimizes the model projection online through the vector g in (53b), selecting the best data-driven prediction at each time step. In contrast, SPC relies on a fixed reduced model computed offline, which cannot adapt to changing operating conditions. Additionally, the Pareto curve shapes differ between the two methods. When the economic weight $\lambda_{economic}$ becomes too large, DeePC's regularization terms ($\lambda_y \|\sigma_y\|^2$ and $\lambda_g \|g\|_p$) compete with the economic objective in the same cost function. This competition causes DeePC to prioritize energy savings over data consistency, resulting in a flatter Pareto front at high comfort violations. SPC, lacking such regularization flexibility, maintains a steeper curve throughout.

This case study demonstrates that DeePC provides superior performance over model-based MPC. It also shows competitive results with SPC under standard operating conditions with stricter comfort criteria. The choice between DeePC and SPC depends on multiple factors. First, consider the priority

between performance optimality and computational cost. If achieving better trade-offs between energy consumption and thermal comfort is the primary goal, DeePC is the preferred choice. Second, consider the available historical data. DeePC demonstrates larger tolerance to limited data information compared to model-based approaches. When high-quality historical data are available but explicit modeling is difficult, DeePC provides a practical alternative that maintains good performance.

2.3 Sensitivity Analysis of DeePC Hyperparameters

Hyper-parameter selection critically influences both economic performance and computational tractability of data-driven predictive control. This subsection analyzes the choice of initial condition estimation order T_{ini}, regularization methodology and weighting, and computational time across varied configurations. The analysis provides practical guidelines for deploying hierarchical DeePC in building thermal management applications.

2.3.1 Initial Condition Estimation Order

The choice of T_{ini} determines the regression order for estimating the initial state of predicted trajectories from past measurements. In the DeePC formulation (III.2), this parameter defines the number of rows in the past input and output matrices U_p and Y_p. The following sensitivity analysis compares ℓ_1 and ℓ_2 regularization methods across different T_{ini} values, providing practical guidelines for controller configuration. The economic indicator I_J evaluates the objective (84) over a week-long simulation, normalized by baseline rule-based control performance. Lower I_J indicates better economic outcomes combining energy savings with comfort maintenance. The computation time indicator I_{time} quantifies the mean wall-clock duration per optimization solve, normalized by the sampling interval $T_s = 15$ minutes according to table III.1.

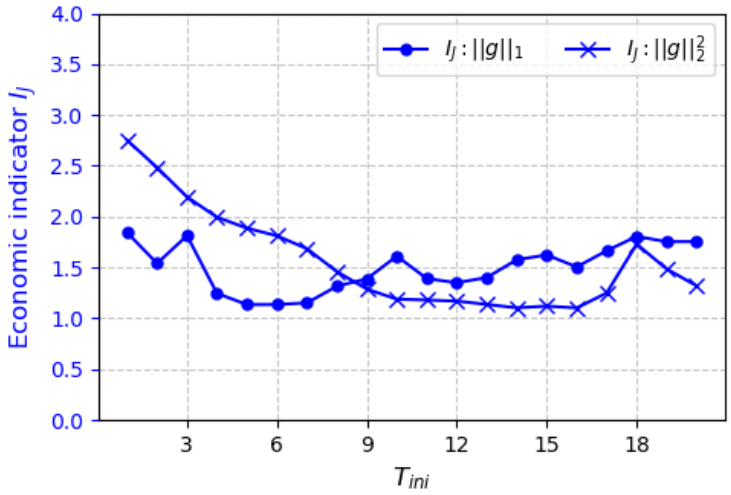

Figure III.5: Economic performance sensitivity to initial state estimation order T_{ini}. Blue circles show ℓ_1 regularization and blue crosses show ℓ_2 regularization. Regularization weights: $\lambda_y = 1$, $\lambda_g = 100$.

Figure III.5 shows the economic performance for both regularization methods. For ℓ_1 regularization (blue circles), the best performance is reached at $T_{ini} \in [4, 7]$, with I_J values near 1.1. Beyond this

range, the performance fluctuates significantly. This behavior occurs because higher order enforces stricter constraints on the initial input sequence u_{ini}. When T_{ini} increases, the past inputs must fit more historical data points, including variations caused by exterior temperature and solar radiation disturbances. This fitting requirement limits the flexibility in choosing optimal power inputs for current predictions. The controller becomes overly constrained by past disturbance patterns that may not be relevant for future control actions, leading to suboptimal economic performance.

In comparison, ℓ_2 regularization (blue crosses) reaches its performance valley at higher $T_{ini} \in [14, 18]$, with I_J decreasing monotonically from 1.12 to approximately 0.94. This difference reflects the properties of each regularization method. The ℓ_1 penalty promotes sparse solutions in vector g, selecting only the most relevant historical trajectory segments. This allows optimal performance with shorter regression windows. In contrast, ℓ_2 regularization encourages smooth solutions that utilize more historical information, requiring larger T_{ini} for optimal state estimation. The ℓ_2 approach aligns better with subspace identification methods that minimize ℓ_2 error, making it more comparable with SPC in terms of model fitting characteristics.

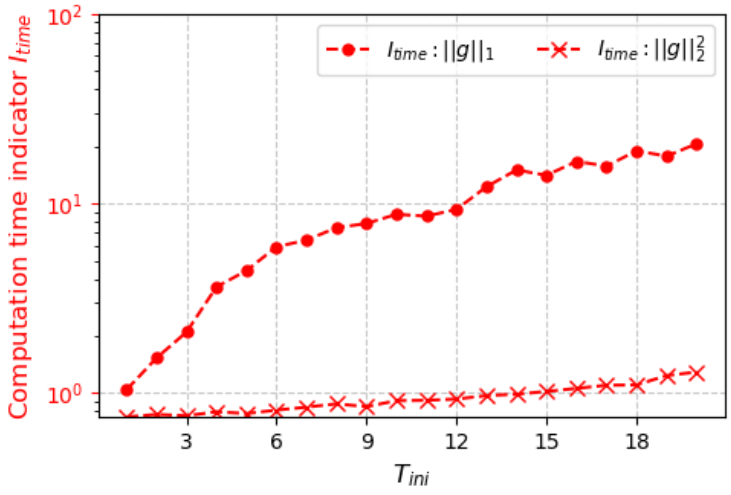

Figure III.6: Computation time sensitivity to initial state estimation order T_{ini}. Red circles show ℓ_1 regularization and red crosses show ℓ_2 regularization on logarithmic scale.

Figure III.6 presents the computational cost comparison. Both regularization methods show exponential growth in computation time as T_{ini} increases, but with different magnitudes. For ℓ_2 regularization (red crosses), I_{time} remains nearly constant around 0.08 across all T_{ini} values, corresponding to approximately 1.2 seconds per solve. This computational stability reflects the quadratic structure of ℓ_2 penalties, which preserve the convex quadratic programming nature of DeePC.

In contrast, ℓ_1 regularization (red circles) incurs substantially higher computational burden. The least squares approximation of ℓ_1 regularization introduces additional decision variables to reformulate the non-smooth absolute value terms. Specifically, expressing $||g||_1 = \sum_i |g_i|$ in standard optimization form requires auxiliary variables and inequality constraints, significantly increasing the problem dimension. At $T_{ini} = 4$, where ℓ_1 achieves its best economic performance, computation time is approximately 3.5$\times$ higher than ℓ_2 alternatives. As T_{ini} increases beyond 10, I_{time} rises from 0.18 to 0.32, potentially challenging real-time implementation. Even though ℓ_2 regularization requires higher

regression order for optimal performance, it still demonstrates a significant advantage in computation time compared to ℓ_1 regularization.

The sensitivity analysis provides practical application guidelines. With higher order T_{ini}, the Hankel matrix size increases proportionally, demanding more memory for optimization. The available memory on the control platform then becomes a limiting factor. Therefore, the choice between ℓ_2 and ℓ_1 regularization depends on hardware configuration and control sampling time. When computational resources are limited or fast sampling is required, ℓ_2 regularization with $T_{ini} \in [14, 18]$ offers robust performance with minimal solver overhead. When memory constraints are severe or training data are sparse, ℓ_1 regularization with $T_{ini} \in [4, 7]$ provides comparable economic outcomes with smaller Hankel matrices, though at the cost of increased computation time per solve. The results emphasize that optimal T_{ini} selection requires balancing economic performance, computational efficiency, and memory availability based on specific deployment constraints.

2.3.2 Regularization Weight Configuration

Regularization weights λ_y and λ_g in equation (III.2) govern the balance between economic objectives, initial state estimation accuracy, and trajectory combination complexity. Insufficient regularization permits overfitting to historical data, yielding predictions that poorly generalize to current operating conditions. Excessive regularization suppresses meaningful trajectory combinations, forcing reliance on trivial solutions that ignore system-specific dynamics embedded in the dataset.

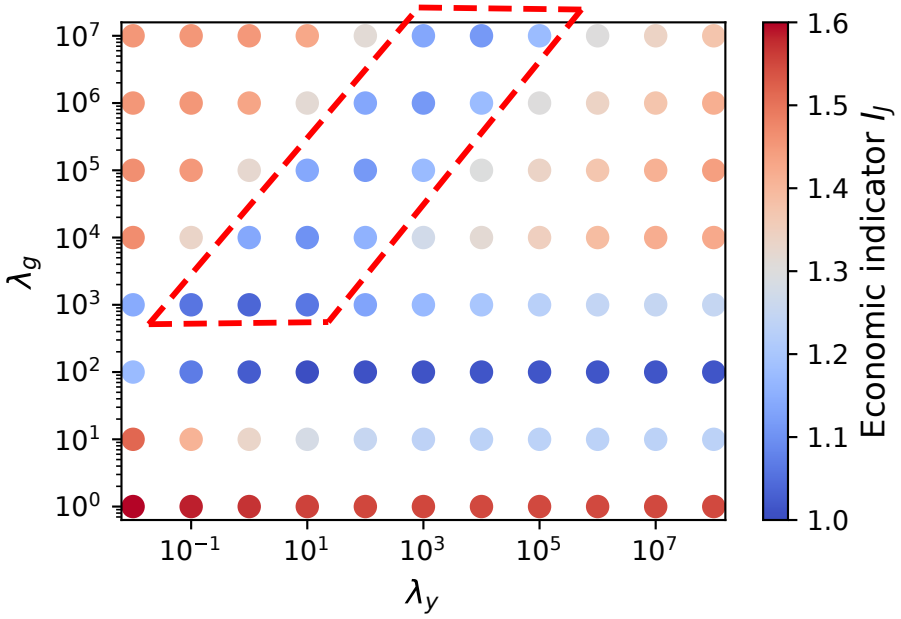

Figure III.7: Economic performance sensitivity to regularization weights λ_y and λ_g. Color intensity indicates normalized economic objective I_J (darker blue represents superior performance). White contours denote iso-performance lines. Red dashed ellipse highlights the approximately optimal parameter region. Diagonal dashed line indicates the parameter relationship corresponding to indirect data-driven predictive control formulation. $T_{ini} = 16$.

Figure III.7 presents economic performance across a two-dimensional grid spanning $\lambda_y \in [10^{-2}, 10^8]$ and $\lambda_g \in [10^0, 10^7]$. Each point represents a week-long closed-loop simulation with

corresponding weight configuration, evaluated using the normalized economic criterion I_J. The heat map coloring encodes performance: darker blue regions indicate superior outcomes (lower I_J), while lighter regions correspond to degraded performance.

For very small $\lambda_g \leq 10^1$, economic performance degrades steeply regardless of λ_y selection, as evidenced by the light-colored left edge of the map. This regime corresponds to essentially unregularized trajectory combination, where the optimization permits arbitrary values in vector g without penalization. Consequently, the constraint (III.2) admits solutions that exactly fit past measurements including noise, while predictions for future horizons become unreliable extrapolations. The resulting control commands exhibit excessive variability, actuating heating equipment erratically in response to spurious patterns.

The best performance is achieved when $\lambda_y \geq 10^0$ and $\lambda_g \approx 10^2$. The unbounded choice of λ_y implies a relatively strict constraint on the state estimation error, allowing the controller to prioritize prediction accuracy while optimizing economic objectives.

In the figure, another region (marked by dashed red line) demonstrates relatively good performance. This region starts from $(10^0, 10^3)$ and slopes upward to the right with a gradient of approximately 1000. This parameter set corresponds to an approximation of bi-level data-driven predictive control, as outlined in [121]. In the bi-level formulation, the initial condition estimation and g regularization objectives are prioritized above the economic objectives. These regularization terms in (53a) are explicitly addressed by solving them as a sub-optimization problem subject to constraint (53b). The resulting solution is then enforced as a hard linear constraint that maps u to g. This approach reduces the number of decision variables since g is no longer jointly optimized with economic objectives. However, this separation prevents the economic objectives from regularizing the underlying data-driven model. The trade-off is computational efficiency at the cost of reduced flexibility in balancing prediction accuracy and economic performance. When the ratio λ_g/λ_y (corresponding to the weighting in the sub-optimization problem) is appropriately chosen, the controller maintains similar performance even if both regularization weights are far higher than the economic objective weights. Similar sloping regions have been observed in other applications [122], suggesting this is a general characteristic of DeePC parameter selection.

When $\lambda_g = 10^2$ and $\lambda_y \geq 10^2$, the initial condition estimation error term becomes dominant in the objective function. In this configuration, the optimization focuses primarily on minimizing σ_y, ensuring that predicted outputs closely match the patterns observed in the historical dataset. This emphasis on fitting past behavior weakens the influence of future trajectory selection based on economic objectives. The input trajectories become heavily constrained by the requirement to match historical state evolution, limiting the controller's ability to adapt to different disturbances or economic incentives. While this configuration ensures prediction consistency with training data, it may sacrifice economic performance when operating conditions deviate from those represented in the Hankel matrix. This observation explains why the performance in this high-weight region, though acceptable, does not match the optimal region where regularization and economic objectives are more balanced.

3 Chapter Conclusion and Outlook

3.1 Conclusion

This chapter investigated hierarchical DeePC for building thermal management, addressing two research questions:

- What distinguishes DeePC from SPC in practice?
- What advantages does DeePC provide for building thermal control?

The key difference lies in when the prediction model is determined. SPC extracts a fixed parametric model offline through subspace projection. DeePC determines the prediction model online by optimizing the trajectory combination vector g at each time step. This allows DeePC to select the Hankel matrix subspace that best approximates current operating conditions. Additionally, the soft constraint σ_y weighted by λ_y controls prediction conservativeness. High λ_y makes the controller prefer trajectories closer to observed dataset behavior. Low λ_y allows exploration beyond strict data consistency, prioritizing economic objectives. The validation confirmed these advantages: DeePC achieved 1°C·hour lower discomfort compared to SPC at equivalent energy consumption.

DeePC addresses a critical challenge in building thermal control: the difficulty of obtaining persistently exciting training data. Thermal experiments are constrained by comfort requirements during occupied periods, limiting temperature variations. Weather disturbances like ambient temperature and solar radiation cannot be manipulated to generate informative excitation patterns. Under such imperfect data conditions, DeePC optimizes the prediction model together with control objectives rather than separating identification and control stages. The regularization terms prevent predictions from diverging from historical observations while maintaining flexibility to achieve economic goals. The hierarchical framework integrating DeePC with VRFT-tuned tracking controllers provides a complete model-free solution that eliminates the modeling bottleneck while maintaining robust performance under realistic data limitations.

From a state-of-the-art perspective, this work extends the comparison between direct and indirect data-driven predictive control to parameter-varying building systems. While recent literature established mathematical equivalence under ideal conditions [120], this chapter demonstrates that practical considerations—regularization design, data quality, and online adaptability—create substantial performance differences.

3.2 Outlook: Adaptive DeePC for Time-Varying Dynamics

From a building control perspective, system thermal conditions change continuously throughout the year. Seasonal variations alter the dominant heat transfer mechanisms and system operating points. Winter operation emphasizes heating capacity and insulation effectiveness. Occupancy patterns evolve with holidays and working schedules. Equipment characteristics drift due to aging and maintenance cycles. These temporal variations mean that a prediction model trained on data from one season may not accurately represent system behavior in another season. Adaptivity is therefore required to recalibrate the prediction model as operating conditions evolve, ensuring that the controller maintains performance across different thermal regimes.

DeePC possesses significant potential for adaptive prediction due to the inherent flexibility of its Hankel matrix formulation. Unlike parametric models that require reidentification procedures when new data become available, DeePC can incorporate new measurements directly into the Hankel matrix structure. Adding recent trajectories simply involves appending new columns to the existing data matrices, without requiring model structure selection or parameter estimation algorithms. This

structural advantage makes DeePC naturally suited for online adaptation. This flexibility suggests that adaptive DeePC could track time-varying dynamics efficiently than methods requiring explicit model reidentification.

However, DeePC faces a unique constraint that distinguishes it from other indirect data-driven approaches. The Hankel matrix dimensions directly determine the optimization problem size and computational complexity. Each column in the Hankel matrix represents a trajectory segment from the historical dataset, and the total number of columns corresponds to the dataset length minus the trajectory horizon. As new data accumulate over extended operation periods, unlimited growth of the Hankel matrix would eventually exceed computational resources and violate real-time control requirements. Therefore, unlike subspace identification methods that can process long datasets by extracting fixed-dimension parametric models, DeePC must carefully manage data selection. The incoming measurements should be strategically chosen to maintain prediction quality while respecting computational constraints. This limitation introduces a trade-off between retaining historical knowledge and incorporating recent observations.

Several aspects must be considered for adaptive DeePC implementation. The next chapter addresses the following questions:

1. What data should be selected to update the Hankel matrix? The new measurements should provide informative content that enriches the behavioral characterization rather than redundant information already captured by existing trajectories.

2. How should new data be incorporated into the Hankel structure? Options include appending new columns while removing old ones, replacing specific trajectory segments, or maintaining a sliding window of recent measurements.

3. How can the consistency of Hankel matrix size be maintained throughout adaptation? The optimization problem formulation assumes fixed matrix dimensions, so any update strategy must preserve the problem structure to ensure solver compatibility.

These questions motivate the investigation of adaptive mechanisms in the next chapter, where different update strategies will be formulated, analyzed, and validated on building thermal management applications.

Chapter IV

Adaptive Data-driven Predictive Control

1 Preliminaries on Adaptive Predictive Control 105
1.1 Motivation and Research question 105
1.2 Literature review on Adaptive MPC 106
2 Adaptive DPC: Classification and Preliminaries 110
2.1 Adaptive Data Update Strategies 110
2.2 Data Quality Assessment 114
3 Adaptive Strategy Validation 117
3.1 Numerical Case Study: RC Parameter Variance 117
3.2 MOSA benchmark validation 128

1 Preliminaries on Adaptive Predictive Control

1.1 Motivation and Research question

Chapter III presented, analyzed, and compared the SPC and DeePC approaches. The advantages and disadvantages of each approach were highlighted, particularly through their application to the MOSA benchmark. However, it is essential to note that these approaches were originally designed under the assumption of a system that remains unchanged over time. This is not the case for the systems we are interested in (building thermal systems), whose models evolve according to operational conditions or environmental factors such as weather.

Extensions of Willems' fundamental lemma to linear parameter-varying (LPV) systems have been proposed in the literature [110]. However, their application to building thermal systems remains underdeveloped. Most existing studies address theoretical aspects. How DeePC maintains prediction accuracy in practical building applications is not well characterized. This gap is significant because building systems continuously experience parameter variations. For example, the wall heat transfer coefficient varies with exterior temperature. Besides, these variations result from weather changes, occupancy patterns, operational disturbances and equipment aging.

The core challenge lies in understanding how parameter variations affect the Hankel data matrix and prediction quality. DeePC operates directly on measured input-output trajectories without explicit model identification. When building dynamics evolve seasonally, historical data in the Hankel matrix becomes increasingly mismatched with current system behavior, potentially degrading prediction accuracy.

This work addresses these gaps by interpreting the DeePC predictor through parameter-varying system theory. We clarify how the Hankel matrix relates to system dynamics when parameters vary

with operating conditions. We explain the role of matrix regularization in maintaining prediction robustness under parameter drift. Based on this understanding, we combine adaptive mechanisms for DeePC: data window strategies, matrix reduction approaches, and forgetting factors. These mechanisms maintain performance as building conditions evolve.

The research question is: how to combine an effective adaptive strategy for DeePC in building thermal control? This study explores it in following sections: section 1.1 reviews the development of Adaptive MPC, which established principles for handling time-varying systems. Section 1.2 examines Willems' lemma under LPV conditions. Section 2 details the core adaptation mechanisms: data window selection, matrix reduction, and forgetting strategies. Section 3 demonstrates these mechanisms through a numerical study with parameter shifts. Section 4 validates the complete adaptive algorithm on MOSA test case.

1.2 Literature review on Adaptive MPC

Adaptive control theory emerged from the foundational work of [123] on self-optimizing control systems, which introduced the principle of automatic parameter adjustment to maintain performance under uncertain conditions. As formalized by [124], adaptive control systems modify their parameters online based on system measurements to achieve desired performance objectives despite parametric variation. The field established two main paradigms: model reference adaptive control (MRAC) that adjust controller parameters to match a reference model, and self-tuning regulators (STR) that combine system identification with control design. Both approaches implement parallel-loop architectures where parameter estimation operates alongside conventional feedback control. These approaches address the challenge of controlling systems with uncertain or time-varying parameters. The theoretical foundation is established for extending adaptive principles to predictive control strategies.

Adaptive MPC was initially introduced because traditional MPC assumes fixed, accurately known plant models. This assumption rarely holds in practice for real industrial systems. The early pioneering works [125, 126] demonstrated the practical necessity of combining online parameter identification with predictive control. The theoretical foundation was later formalized by [127, 128] through Generalized Predictive Control (GPC). GPC explicitly incorporated recursive parameter estimation within the predictive control framework. This paradigm addresses a challenge: system dynamics evolve continuously due to equipment degradation, changing operating conditions, and modeling uncertainties.

The field experienced significant theoretical advances during the 1990s and 2000s. A comprehensive theoretical basis for adaptive MPC emerged during this period, with issues of feasibility, stability, and performance becoming largely understood for linear systems [129, 130]. The 2010s witnessed growing interest in robust adaptive approaches that address constraint satisfaction under parameter uncertainty. These developments introduced techniques such as tube-based MPC and set-membership identification to handle parametric uncertainties while maintaining feasibility guarantees [131, 132].

Recently, the integration of data-driven techniques has opened new research directions [133]. Advanced uncertainty quantification techniques using Bayesian neural networks and meta-learning approaches have emerged to better capture and adapt to time-varying uncertainties [134, 135]. Deep learning integration has progressed toward real-time implementations, with neural network-based adaptive MPC frameworks being successfully deployed on embedded systems and autonomous platforms [136, 137].

Benefiting from convenient updating mechanisms, adaptive frameworks is introduced in DPC. Adaptive Subspace Predictive Control (ASPC) incorporates recursive subspace identification that continuously update system models using sliding data windows [30, 138]. ASPC has demonstrated success in building applications. It enables real-time tracking of thermal dynamics changes due to seasonal variations and occupancy patterns [139]. Adaptive direct data-driven predictive control (ADDPC), addressed similar challenges by continuously refreshing the data matrices used in direct data-driven predictive control (DDPC) formulations without explicit model identification [104, 102]. Both parametric and non-parametric data-driven methods have developed to provide robust solutions for time-varying systems[45].

1.2.1 Predictor interpretation for LPV systems

Consider an LPV system where state-space matrices depend on a scheduling parameter $\theta[k]$. For building thermal systems, $\theta[k]$ represents thermal dynamic coefficients (like $\boldsymbol{\theta_{phi}}$ in equation (19)) that cause system behavior to shift between operating regimes. For simplicity of notation, we denote the parameter-dependent matrices as $(A_k, B_k, C_k, D_k) := (A(\theta[k]), B(\theta[k]), C(\theta[k]), D(\theta[k]))$ throughout this section. The considered LPV system is:

$$\begin{aligned} x[k+1] &= A_k x[k] + B_k u[k], \\ y[k] &= C_k x[k] + D_k u[k]. \end{aligned} \tag{86}$$

Following the definition of the trajectory simulation matrix Γ for LTI systems in equation (59), the historical dataset allows us to construct matrices for LPV systems. The Hankel matrices satisfy:

$$\begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} = \mathbf{\Gamma}^{data} \begin{bmatrix} X_p \\ U_p \\ U_f \end{bmatrix}, \tag{87}$$

where $\mathbf{\Gamma}^{data}$ is a collection of trajectory simulation matrices indexed by time instants in the historical data:

$$\mathbf{\Gamma}^{data} = [\Gamma^{data}[0], \Gamma^{data}[1], \ldots, \Gamma^{data}[T_d - T_{ini} - N + 1]], \tag{88}$$

The length of vector corresponds to the column number of Hankel matrices. In equation (88), (U_p, Y_p, U_f, Y_f) are Hankel matrices. X_p is system initial condition, as defined in equation (55). Each pair of data columns in equation (88) is bridged by a different $\Gamma^{data}[k]$:

$$\Gamma^{data}[k] = \begin{bmatrix} \mathbf{0} & I_{T_{ini}m} & \mathbf{0} \\ \bar{\mathcal{O}}^{data}_{T_{ini}}[k] & \bar{\mathcal{T}}^{data}_{T_{ini}}[k] & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & I_{Nm} \\ \bar{\mathcal{O}}^{data}_{N}[k]\mathcal{A}_{k+T_{ini}-1:k} & \mathbf{0} & \bar{\mathcal{T}}^{data}_{N}[k] \end{bmatrix}, \tag{89a}$$

$$\text{with} \quad \mathcal{A}_{i:j} := A_i A_{i-1} \cdots A_j, \quad \forall i \geq j. \tag{89b}$$

For LPV systems, the trajectory simulation matrix is constructed using adapted forms of Toeplitz and observability matrices, denoted with a bar as ($\bar{\mathcal{T}}^{data}_t[k]$ and $\bar{\mathcal{O}}^{data}_t[k]$). At time instant k, these

matrices are:

$$\bar{\mathcal{T}}_t^{data}[k] := \begin{bmatrix} D_k & 0 & 0 & \cdots & 0 \\ C_{k+1}B_k & D_{k+1} & 0 & \cdots & 0 \\ C_{k+2}\mathcal{A}_{k+1:k}B_k & C_{k+2}B_{k+1} & D_{k+2} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ C_{k+t-1}\mathcal{A}_{k+t-2:k}B_k & C_{k+t-1}\mathcal{A}_{k+t-2:k+1}B_{k+1} & C_{k+t-1}B_{k+t-2} & \cdots & D_{k+t-1} \end{bmatrix}, \tag{90a}$$

$$\bar{\mathcal{O}}_t^{data}[k] := \begin{bmatrix} C_k \\ C_{k+1}\mathcal{A}_{k:k} \\ C_{k+2}\mathcal{A}_{k+1:k} \\ \vdots \\ C_{k+t-1}\mathcal{A}_{k+t-2:k} \end{bmatrix}. \tag{90b}$$

where t is the order of matrix, similar to definition of toeplitz and observability matrices in equation (45).

In equation (87), $\mathbf{\Gamma}^{data}$ represents a collection of system behavior snapshots at different time instants in the historical data. Each $\Gamma^{data}[k]$ captures the system dynamics under the parameter values experienced at instant k, separating the influence of initial conditions, past inputs, and future inputs.

On the right side of the prediction constraint in equation (53b), the online regression for future trajectory can be decomposed similarly as:

$$\begin{bmatrix} \mathbf{u}_{ini} \\ \mathbf{y}_{ini} \\ \mathbf{u} \\ \mathbf{y} \end{bmatrix} = \Gamma^{online}[k_0] \begin{bmatrix} \mathbf{x}[k_0] \\ \mathbf{u}_{ini} \\ \mathbf{u} \end{bmatrix}, \tag{91}$$

where k_0 denotes the time instant of the current measurement and decision. The matrix $\Gamma^{online}[k_0]$ is defined in the same structure as $\Gamma^{data}[k]$:

$$\Gamma^{online}[k_0] = \begin{bmatrix} \mathbf{0} & I_{T_{ini}m} & \mathbf{0} \\ \bar{\mathcal{O}}_{T_{ini}}^{online}[k_0] & \bar{\mathcal{T}}_{T_{ini}}^{online}[k_0] & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & I_{Nm} \\ \bar{\mathcal{O}}_N^{online}[k_0]\mathcal{A}_{k_0+T_{ini}-1:k_0} & \mathbf{0} & \bar{\mathcal{T}}_N^{online}[k_0] \end{bmatrix}, \tag{92a}$$

with the pseudo-Toeplitz and observability matrices constructed from the parameter values at the prediction time instants, following the same structure as equations (90a) and (90b) with time indices relative to k_0.

The matrix $\Gamma^{online}[k_0]$ represents the system behavior at the current time instant k_0, enabling prediction of future outputs and states based on initial conditions and planned inputs.

In the LTI case from Chapter III (with θ constant as in equation (19)), the trajectory combination vector g is optimized to minimize the state estimation error σ_y:

$$\begin{aligned} \sigma_y &= Y_p g - \mathbf{y}_{ini} && (93a)\\ &= (\mathcal{T}_{T_{ini}} U_p + \mathcal{O}_{T_{ini}} X_p) g - \mathbf{y}_{ini} && (93b)\\ &= (\mathcal{T}_{T_{ini}} U_p + \mathcal{O}_{T_{ini}} X_p) g - (\mathcal{T}_{T_{ini}} \mathbf{u}_{ini} + \mathcal{O}_{T_{ini}} \mathbf{x}[k_0]) && (93c)\\ &= \mathcal{T}_{T_{ini}} (U_p g - \mathbf{u}_{ini}) + \mathcal{O}_{T_{ini}} (X_p g - \mathbf{x}[k_0]). && (93d) \end{aligned}$$

Minimizing σ_y serves two purposes: it combines g to satisfy $U_p g \approx \mathbf{u}_{ini}$ (matching the current input trajectory history) while simultaneously satisfying $X_p g \approx \mathbf{x}[k_0]$ (estimating the current initial state).

A key difference between LTI and LPV systems lies in how the output Hankel matrix Y_p is composed. In the LTI case, Y_p is constructed from a single constant Toeplitz and observability matrix pair. In the LPV case, Y_p is constructed from the collection $\mathbf{\Gamma}^{data}$, which contains a library of pseudo-Toeplitz and pseudo-observability matrices corresponding to different time instants in the historical data.

Define the library of pseudo-Toeplitz and observability matrices extracted from $\mathbf{\Gamma}^{data}$ as:

$$\bar{\boldsymbol{\mathcal{T}}}^{data} = [\bar{\mathcal{T}}^{data}_{T_{ini}}[0], \bar{\mathcal{T}}^{data}_{T_{ini}}[1], \dots, \bar{\mathcal{T}}^{data}_{T_{ini}}[T_d - T_{ini} - N + 1]], \tag{94}$$

$$\bar{\boldsymbol{\mathcal{O}}}^{data} = [\bar{\mathcal{O}}^{data}_{T_{ini}}[0], \bar{\mathcal{O}}^{data}_{T_{ini}}[1], \dots, \bar{\mathcal{O}}^{data}_{T_{ini}}[T_d - T_{ini} - N + 1]], \tag{95}$$

where each pair $(\bar{\mathcal{T}}^{data}_{T_{ini}}[k], \bar{\mathcal{O}}^{data}_{T_{ini}}[k])$ forms part of the matrix $\Gamma^{data}[k]$.

The output Hankel matrix is then:

$$Y_p = \bar{\boldsymbol{\mathcal{T}}}^{data} U_p + \bar{\boldsymbol{\mathcal{O}}}^{data} X_p, \tag{96}$$

where the library matrices are applied to the Hankel data columns through the weighting vector g.

The state estimation error for LPV systems is:

$$\begin{aligned}
\sigma_y &= Y_p g - \mathbf{y}_{ini} && \text{(97a)}\\
&= (\bar{\boldsymbol{\mathcal{T}}}^{data} U_p + \bar{\boldsymbol{\mathcal{O}}}^{data} X_p) g - \mathbf{y}_{ini} && \text{(97b)}\\
&= (\bar{\boldsymbol{\mathcal{T}}}^{data} U_p + \bar{\boldsymbol{\mathcal{O}}}^{data} X_p) g - (\bar{\mathcal{T}}^{online}_{T_{ini}}[k_0]\mathbf{u}_{ini} + \bar{\mathcal{O}}^{online}_{T_{ini}}[k_0]\mathbf{x}[k_0]) && \text{(97c)}\\
&= \bar{\boldsymbol{\mathcal{T}}}^{data}(U_p g - \mathbf{u}_{ini}) + \bar{\boldsymbol{\mathcal{O}}}^{data}(X_p g - \mathbf{x}[k_0]) && \text{(97d)}\\
&\quad + (\bar{\boldsymbol{\mathcal{T}}}^{data} - \bar{\mathcal{T}}^{online}_{T_{ini}}[k_0])\mathbf{u}_{ini} + (\bar{\boldsymbol{\mathcal{O}}}^{data} - \bar{\mathcal{O}}^{online}_{T_{ini}}[k_0])\mathbf{x}[k_0]. && \text{(97e)}
\end{aligned}$$

The two terms in (97d) represent trajectory matching error and initial state estimation error, analogous to the LTI case.

The two terms in (97e) arise from parameter variations and represent model mismatch specific to LPV systems. The term $(\bar{\boldsymbol{\mathcal{T}}}^{data} - \bar{\mathcal{T}}^{online}_{T_{ini}}[k_0])$ quantifies the discrepancy between the library of historical pseudo-Toeplitz matrices and the pseudo-Toeplitz matrix of the current system at time k_0. This mismatch occurs because the parameter values in the historical data differ from those at the current operating point. Similarly, $(\bar{\boldsymbol{\mathcal{O}}}^{data} - \bar{\mathcal{O}}^{online}_{T_{ini}}[k_0])$ captures the difference in observability due to parameter variations, affecting how initial state influences the output trajectory.

This reveals a challenge for the Willems' fundamental lemma-based predictor: when the current parameter $\theta[k]$ differs from the parameter values in the historical datase, the optimization of g will be biased by the residue of (97e), disturbing the composition of the trajectory and initial condition estimation.

Comparing LPV with the LTI case, the mismatch between dataset behavior and real-time behavior becomes an additional error terms in σ_y minimization. Two key insights is listed for DDPC applied to LPV systems:

1. The Hankel matrix should record system behavior similar to closed-loop operation to reduce initial state estimation errors. When dataset parameter conditions differ substantially from current conditions, prediction accuracy degrades.

2. The decision variable g only rely on the underlying model terms in (97d), which is encoded by the pseudo trajectory dynamic matrices $\bar{\boldsymbol{\mathcal{T}}}^{data}$ and $\bar{\boldsymbol{\mathcal{O}}}^{data}$ constructed from historical data. The bias is not related to decision variable g.

This observation aligns with perspectives in [115, 140]. It suggests two distinct approaches to enhance DeePC performance under parameter variations. The first approach updates the Hankel data matrix continuously, replacing outdated trajectories with recent measurements representative of current system behavior. This reduces the mismatch terms in Equation (97) by keeping Γ^{data} aligned with $\Gamma[k]$. The second approach estimates and compensates for the mismatch error terms before optimization, treating them as known disturbances that can be predicted and rejected. This chapter focuses on the first approach through adaptive data update mechanisms. The second approach remains an open direction for future investigation.

2 Adaptive DPC: Classification and Preliminaries

2.1 Adaptive Data Update Strategies

Although Adaptive SPC (ASPC) has a longer history than Adaptive DDPC (ADDPC), both methodologies exhibit a duality in their development within the adaptive control framework. The distinction lies in their treatment of prediction models: ASPC performs indirect data-driven prediction by deriving a parametric state-space model through recursive estimation, while ADDPC bypasses explicit model identification and directly uses the data matrix column space for prediction. Despite this methodological difference, both approaches share a common theoretical foundation in exploiting the column space of data matrices for system prediction. This shared nature directs both adaptive schemes toward the same focus: how to update data matrices effectively.

The adaptive DeePC framework processes new measurement pairs $\{u(k), y(k)\}$ at each time step k to update the Hankel matrix for the next control iteration. This online construction maintains alignment with evolving system dynamics without requiring offline reidentification procedures. The core challenge is managing the trade-off between maintaining rich information content and controlling computational complexity.

Three primary mechanisms address this challenge. Each offers distinct advantages and trade-offs in memory management, computational cost, and adaptation responsiveness.

2.1.1 Sliding Window Methods

Sliding window methods maintain fixed matrix dimensions by removing older data columns and appending new measurements. This approach is computationally straightforward and preserves operational simplicity.

Let $\mathscr{H}_k := \mathscr{H}_{T_{ini}+N}(\mathbf{u}_d^k, \mathbf{y}_d^k)$ denote the Hankel matrix at time step k, constructed from historical input-output data sequences $\{\mathbf{u}_d^k, \mathbf{y}_d^k\}$ of length T_k. At each time step, a new data column candidate $\mathbf{c}_k$ becomes available:

$$\mathbf{c}_k = \begin{bmatrix} u[k+1-T_{ini}-N] \\ \vdots \\ u[k] \\ y[k+1-T_{ini}-N] \\ \vdots \\ y[k] \end{bmatrix} \in \mathbb{R}^{(n_u+n_y)(T_{ini}+N)} \tag{98}$$

The Hankel matrix at time k with sliding window length N_{win} is:

$$\mathscr{H}_k = \begin{bmatrix} \mathbf{c}_{k-N_{\text{win}}+1} & \mathbf{c}_{k-N_{\text{win}}+2} & \cdots & \mathbf{c}_k \end{bmatrix} \tag{99}$$

When new measurement arrives at time k, the update removes the oldest column and appends the newest from the Hankel in previous time step:

$$\mathscr{H}_{k-1} = \begin{bmatrix} \mathbf{c}_{k-N_{\text{win}}} & \mathbf{c}_{k-N_{\text{win}}+1} & \cdots & \mathbf{c}_k - 1 \end{bmatrix} \tag{100}$$

The matrix dimensions remain constant at $(n_u + n_y)(T_{\text{ini}} + N) \times N_{\text{win}}$, providing constant computational cost for subsequent optimization. This fixed structure is particularly valuable for real-time control where computational load must be predictable.

For ASPC, sliding window methods use rolling-window techniques in recursive subspace identification algorithms [141, 142]. For ADDPC, sliding window methods retain informational richness within the Hankel matrix structure [112, 143].

The two approaches exhibit differences in their flexibility. Subspace identification algorithms can process relatively large datasets because the identification procedure inherently summarizes features across extensive data collections, enabling ASPC to utilize wider sliding windows with lower computational cost. In contrast, ADDPC faces more restrictive limitations since the data matrix directly participates in online optimization procedures. The window width N_{win} determines the Hankel matrix size, which significantly affects computational load. This necessitates smaller window sizes to maintain feasible optimization times in real-time applications.

To address the trade-off between data richness and computational efficiency, column subset selection (CSS) algorithms evaluate incoming data columns against quality criteria before incorporation [144]. Rather than automatically accepting all new measurements, this selective approach enables strategic data management while maintaining matrix size constraints.

Several implementations demonstrate these principles:

- **Residual-based validation**: Quadratic ARX model residuals evaluate new data quality, triggering dataset resampling via column subset selection when residuals exceed predefined tolerances [144].
- **Schur complement method**: The persistent excitation (PE) condition is reformulated as a linear matrix inequality (LMI), transforming the problem into a semi-definite program (SDP). While this rigorously guarantees the PE condition, it significantly increases computational complexity [145].
- **Leverage score sampling**: Column selection uses leverage scores based on L2 norm distributions to guide randomized sampling for matrix reconstruction when data quality deteriorates [146].

2.1.2 Recursive Reduction Methods

Recursive reduction methods allow the Hankel matrix to grow by appending new measurements, then compress the expanded matrix through dimensionality reduction. Rather than maintaining fixed dimensions or implementing fixed-size windows, these approaches strategically compress accumulated information to retain relevant system dynamics while controlling computational complexity.

Starting from Hankel matrix at time $k-1$:

$$\mathscr{H}_{k-1} = \begin{bmatrix} \mathbf{c}_1 & \mathbf{c}_2 & \cdots & \mathbf{c}_{k-1} \end{bmatrix} \tag{101}$$

At time k, a new column $\mathbf{c}_k$ is appended, creating an expanded matrix:

$$\bar{\mathscr{H}_k} = \begin{bmatrix} \mathbf{c}_1 & \mathbf{c}_2 & \cdots & \mathbf{c}_{k-1} & \mathbf{c}_k \end{bmatrix} = \begin{bmatrix} \mathscr{H}_{k-1} & \mathbf{c}_k \end{bmatrix} \tag{102}$$

The compression uses singular value decomposition (SVD) to extract the dominant subspace:

$$\bar{\mathscr{H}_k} = U_k \Sigma_k V_k^\top \tag{103}$$

where $U_k \in \mathbb{R}^{(n_u+n_y)(T_{\text{ini}}+N)\times r_k}$ contains left singular vectors, $\Sigma_k \in \mathbb{R}^{r_k \times r_k}$ contains singular values in descending order, and r_k represents the numerical rank.

The updated matrix retains only the left singular vectors and singular values:

$$\mathscr{H}_k = U_k \Sigma_k \tag{104}$$

This compressed representation maintains fixed row dimensions $(n_u+n_y)(T_{\text{ini}}+N)$ while column count remains at $r_k \leq (n_u+n_y)(T_{\text{ini}}+N)$. The matrix naturally reaches a square structure, after which dimensions stabilize. The optimization remains analytically equivalent, but more efficient on memory usage [147].

For ASPC, dimensionality reduction occurs naturally through online subspace identification, which inherently compresses information by extracting low-dimensional state-space representations. For ADDPC, SVD-based compression removes null-space components from the right companion matrix, reconstructing a square data matrix that maintains stable computational costs regardless of original dataset size [148, 149, 150].

The computational cost of full SVD at each time step can become prohibitive. Full recomputation requires $\mathcal{O}((n_u+n_y)(T_{\text{ini}}+N)\cdot k^2)$ operations, growing quadratically with accumulated data. Fast-SVD algorithms enable efficient recursive updates by exploiting the rank-one modification structure of column appending [147].

Given previous decomposition $\mathscr{H}_{k-1} = U_{k-1}\Sigma_{k-1}$ and new column $\mathbf{c}_k$, the algorithm first computes:

$$\mathbf{p} = U_{k-1}^\top \mathbf{c}_k, \quad \mathbf{q} = \mathbf{c}_k - U_{k-1}\mathbf{p} \tag{105}$$

The vector $\mathbf{p}$ represents the projection of the new column $\mathbf{c}_k$ onto the subspace spanned by the left singular vectors U_{k-1}. It captures the component of $\mathbf{c}_k$ that lies within the existing low-dimensional subspace. The vector $\mathbf{q}$ represents the orthogonal residual—the part of $\mathbf{c}_k$ orthogonal to the existing subspace. When $\mathbf{q}$ is small, the new data adds little new information; when $\mathbf{q}$ is large, the new data significantly expands the subspace.

Then constructs an augmented matrix and computes its small SVD:

$$\tilde{\Sigma}_k = \begin{bmatrix} \Sigma_{k-1} & \mathbf{p} \\ \mathbf{0}^\top & \|\mathbf{q}\|_2 \end{bmatrix}, \quad \tilde{\Sigma}_k = \tilde{U}_k \Sigma_k \tilde{V}_k^\top \tag{106}$$

The updated singular vectors become:

$$U_k = \begin{bmatrix} U_{k-1} & \mathbf{q}/\|\mathbf{q}\|_2 \end{bmatrix} \tilde{U}_k \tag{107}$$

This reduces computational complexity to $\mathcal{O}((n_u + n_y)(T_{\text{ini}} + N) \cdot r_k^2)$, where r_k is typically much smaller than accumulated data count k.

The SVD-based compression naturally supports weighting strategies through diagonal scaling of columns before decomposition. By applying monotonically decreasing weights to fade older measurements before SVD computation, the compressed representation automatically prioritizes recent dynamics. This property bridges to forgetting mechanisms, where temporal weighting enhances adaptation responsiveness.

2.1.3 Forgetting Mechanisms

Forgetting mechanisms implement temporal weighting strategies that assign exponentially decaying weights to historical data while preserving recent information dominance. This provides continuous adaptation, avoiding the discrete information loss characteristic of window removal.

Both sliding window and recursive reduction methods treat retained data equally regardless of age. For time-varying systems, this becomes problematic: older measurements reflect past dynamics no longer representative of current behavior. Forgetting mechanisms prioritize recent information through exponential decay.

The classical approach applies uniform temporal decay via forgetting factors $\gamma \in (0, 1]$. In recursive reduction [151], the updated singular vectors are weighted as:

$$U_k = \begin{bmatrix} U_{k-1} & \gamma\mathbf{q}/\|\mathbf{q}\|_2 \end{bmatrix} \tilde{U}_k \tag{108}$$

Recent columns maintain near-unity weights while distant columns decay exponentially toward zero.

The effective memory length is characterized by the threshold where weights become negligible. Defining significance threshold $\epsilon_w = 0.05$, the effective window size becomes:

$$N_{\text{eff}} = \left\lceil \frac{\ln \epsilon_w}{\ln \gamma} \right\rceil \tag{109}$$

For $\gamma = 0.99$ and $\epsilon_w = 0.05$, this yields $N_{\text{eff}} \approx 300$ samples. With hourly sampling in building control, this corresponds to approximately 12.5 days of effective memory, allowing the controller to track slow thermal dynamics variations while progressively discarding outdated information. Reducing the forgetting factor produces shorter memory: $\gamma = 0.95$ yields $N_{\text{eff}} \approx 60$ samples.

First, the forgetting factor γ offers continuous tuning of memory length, unlike sliding windows constrained to discrete multiples of window size. Second, exponential decay creates smooth transitions when data loses relevance, avoiding discontinuous jumps from window removal. Third, the mechanism integrates naturally with both paradigms through simple diagonal multiplication, requiring minimal computational overhead.

The DeePC requirement for persistent excitation constrains minimum column count:

$$T_{\min} = (n_u + n_y)(T_{\text{ini}} + N) \tag{110}$$

Sliding window methods require window size $N_{\text{win}} \geq T_{\min}$ at all times, strictly limiting memory to discrete multiples. In contrast, forgetting mechanisms decouple memory from matrix dimensions. The Hankel matrix maintains column count at least $T_{\min}$ for PE satisfaction, while effective memory N_{eff} is independently tuned. When $N_{\text{eff}} < T_k$, older columns remain present but contribute negligibly through exponentially small weights. This flexibility enables shorter adaptation timescales without violating PE conditions.

Sliding windows record exact measured trajectory segments, preserving raw column space structure and maintaining direct data-driven prediction where forecasts combine measured trajectories. Recursive reduction with forgetting constructs an optimal basis representing essential information across weighted measurements. The resulting column space provides equivalent predictive capability through orthogonalized, compressed representation rather than direct trajectory storage.

Uniform forgetting can induce estimator windup in poorly excited parameter directions, leading to numerical instability [152]. Directional forgetting applies decay selectively along directions corresponding to new information arrival [153]. Advanced matrix forgetting factor approaches combine variable-rate and variable-direction forgetting through adaptive forgetting matrices $\Lambda(k)$ responding to local excitation characteristics [154]. Recent implementations integrate forgetting with data quality assessment, adjusting forgetting rates based on residual validation and PE monitoring to maintain performance across seasonal transitions while preventing degradation from corrupted measurements [144, 155].

2.2 Data Quality Assessment

The adaptive update mechanism for the Hankel matrix requires evaluation of incoming data quality to ensure persistent excitation and model accuracy. This subsection introduces three complementary data quality indicators that collectively assess whether new measurements should be incorporated into the Hankel matrix.

The decision to incorporate $\mathbf{c}_k$ into $\mathscr{H}_k$ depends on three quality indicators: I_{match}, I_{info}, and I_{oper}.

Indicator 1: Model Matching (I_{match})

The model matching indicator quantifies how well the current model predicts the system behavior. At time k, the one-step-ahead prediction error is:

$$e_{\text{pred}}[k] = \|y[k] - \hat{y}[k|k-1]\|_2^2 \tag{111}$$

where $\hat{y}[k|k-1]$ is the predicted output at time k based on information available at $k-1$, obtained from the DeePC optimization. Specifically, from the optimal solution $g^*[k-1]$ computed at time $k-1$:

$$\hat{y}[k|k-1] = [Y_f g^*[k-1]]_{1:n_y} \tag{112}$$

where $[\cdot]_{1:n_y}$ extracts the first n_y elements corresponding to the first prediction step.

The model matching indicator over a sliding window of length N_w is:

$$I_{\text{match}}[k] = \exp\left(-\frac{1}{N_w \bar{e}_{\text{pred}}} \sum_{j=k-N_w+1}^{k} e_{\text{pred}}[j]\right) \tag{113}$$

where $\bar{e}_{\text{pred}}$ is a reference error threshold. Higher I_{match} (approaching 1) indicates better model accuracy. When $I_{\text{match}}[k]$ falls below a threshold $\bar{I}_{\text{match}}$, it suggests model degradation, potentially warranting data incorporation to refresh the Hankel matrix.

The window length N_w should balance noise filtering with responsiveness. Values between 5 and 20 time steps provide effective averaging. Buildings with slower thermal dynamics benefit from longer windows ($N_w = 15 - 20$), while lightweight buildings require shorter windows ($N_w = 5 - 10$) to detect model mismatch quickly. The threshold $\bar{e}_{\text{pred}}$ should be calibrated relative to typical squared temperature error: conservative settings use $\bar{e}_{\text{pred}} = 0.1 - 0.3$ K^2 for comfort-critical applications, while energy-focused applications tolerate $\bar{e}_{\text{pred}} = 0.5 - 1.0$ K^2. The acceptance threshold is typically $\bar{I}_{\text{match}} = 0.3 - 0.5$.

Indicator 2: Data Informativity (I_{info})

The data informativity indicator evaluates whether the augmented Hankel matrix $\tilde{\mathscr{H}}_k = [\mathscr{H}_{k-1}\ \mathbf{c}_k]$ satisfies the PE condition necessary for reliable prediction. Following the framework in [105], we assess informativity through the rank of the input sub-matrix.

Let $\mathscr{H}^u_{k-1}$ denote the input portion of the Hankel matrix:

$$\mathscr{H}^u_{k-1} = \mathscr{H}_{T_{ini}+N}(\mathbf{u}_d^{k-1}) \in \mathbb{R}^{n_u(T_{ini}+N)\times T_{k-1}} \tag{114}$$

The informativity condition requires:

$$\text{rank}(\mathscr{H}^u_{k-1}) = n_u(T_{ini} + N) \tag{115}$$

We compute the numerical rank using singular value decomposition:

$$\mathscr{H}^u_{k-1} = U_{k-1}\Sigma_{k-1}V_{k-1}^\top, \quad \Sigma_{k-1} = \text{diag}(\sigma_1^{[k-1]}, \ldots, \sigma_{r_{k-1}}^{[k-1]}) \tag{116}$$

where $\sigma_1^{[k-1]} \geq \sigma_2^{[k-1]} \geq \cdots \geq \sigma_{r_{k-1}}^{[k-1]} > 0$ are the singular values. The numerical rank is:

$$r_{\text{num}}[k-1] = \sum_{i=1}^{r_{k-1}} \mathbb{1}\left\{\sigma_i^{[k-1]} > \epsilon_{\text{tol}}\right\} \tag{117}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function that returns 1 if the condition inside is true, and 0 otherwise. Tolerance ϵ_{tol} is either manually defined, or chosen according to a percentage of the total singular value:

$$r_{\text{num}}[k-1] = \sum_{i=1}^{r_{k-1}} \mathbb{1}\left\{\sigma_i^{[k-1]} > \eta \sum_{j=0}^{r_{k-1}} \sigma_j^{[k-1]}\right\} \tag{118}$$

where η is the ignored percentage of singular value.

The informativity indicator is defined as:

$$I_{\text{info}}[k-1] = \frac{r_{\text{num}}[k-1]}{n_u(T_{ini}+N)} \tag{119}$$

For the candidate augmented Hankel $\tilde{\mathscr{H}}_k$:

$$I_{\text{info}}^{\text{aug}}[k] = \frac{r_{\text{num}}(\tilde{\mathscr{H}}_k^u)}{n_u(T_{ini}+N)} \tag{120}$$

Accept $\mathbf{c}_k$ if:

$$I_{\text{info}}^{\text{aug}}[k] \geq \epsilon_{\text{info}} \quad \text{and} \quad I_{\text{info}}^{\text{aug}}[k] \geq I_{\text{info}}[k-1] \tag{121}$$

where $\epsilon_{\text{info}} \in (0, 1]$ is the informativity threshold (typically $\epsilon_{\text{info}} \geq 0.95$).

Values approaching unity ($\epsilon_{\text{info}} = 0.95 - 0.99$) enforce near-full-rank conditions, ensuring robust numerical conditioning but potentially rejecting marginally exciting data. Safety-critical applications warrant higher thresholds. Moderately relaxed thresholds ($\epsilon_{\text{info}} = 0.90 - 0.95$) balance robustness with adaptation flexibility. Excessively low thresholds risk accepting poorly conditioned data that degrades optimization performance.

Indicator 3: Operational Relevance (I_{oper})

For building thermal control systems, the primary control objective is to maintain occupant comfort during service hours. Data collected outside operational periods (e.g., during nights or weekends when buildings are unoccupied) has limited relevance for improving control performance during actual service times.

Define the operational relevance indicator as:

$$I_{\text{oper}}[k] = \begin{cases} 1, & \text{if } t[k] \in \mathbb{T}_{\text{service}} \\ \alpha_{\text{off}}, & \text{otherwise} \end{cases} \tag{122}$$

where $\mathbb{T}_{\text{service}}$ denotes the set of service hours (e.g., weekdays 7:00-19:00), and $0 \leq \alpha_{\text{off}} < 1$ is a reduced weight for off-service data (typically $\alpha_{\text{off}} = 0.1$ or 0).

This indicator ensures that the Hankel matrix prioritizes data from operationally relevant periods, improving the model's predictive accuracy when it matters most. Additionally, for systems with strong diurnal patterns, this prevents the dataset from being dominated by low-activity nighttime data that provides limited information about daytime operation.

The weight α_{off} determines how much off-hours data contributes to the matrix. Setting $\alpha_{\text{off}} = 0$ completely excludes nighttime data, appropriate when day/night dynamics differ substantially. Setting $\alpha_{\text{off}} = 0.1 - 0.3$ retains minimal nighttime information to maintain matrix conditioning while strongly prioritizing service-hour data. The service hours $\mathbb{T}_{\text{service}}$ should match actual occupancy patterns, typically 7:00-19:00 for office buildings.

Unified Decision Logic

The three indicators are combined to determine data acceptance. The logic is expressed in Algorithm 1:

Algorithm 1: Data Quality Assessment for Adaptive Hankel Update

Input: Hankel matrix $\mathscr{H}_{k-1}$, new column $\mathbf{c}_k$, thresholds $\{\bar{I}_{\text{match}}, \epsilon_{\text{info}}, \epsilon_{\text{oper}}\}$
Output: Decision $\in \{\text{Accept}, \text{Reject}\}$

Compute $I_{\text{match}}[k]$ using equation (113)
Compute $I_{\text{info}}^{\text{aug}}[k]$ using equation (119)
Compute $I_{\text{oper}}[k]$ using equation (122)
if $\underline{I_{\text{match}}[k] < \bar{I}_{\text{match}}}$ **then**
    `// Model matching is poor: model degradation detected`
    if $\underline{(I_{\text{info}}^{\text{aug}}[k] \geq \epsilon_{\text{info}}) \wedge (I_{\text{oper}}[k] > \epsilon_{\text{oper}})}$ **then**
        `// New data is informative and operationally relevant: refresh needed`
        return Accept $\mathbf{c}_k$
    else
        `// New data insufficient quality to remedy degradation`
        return Reject $\mathbf{c}_k$
    end
else
    `// Model matching is adequate: check if new data improves matrix`
    if $\underline{(I_{\text{info}}^{\text{aug}}[k] > I_{\text{info}}[k-1]) \wedge (I_{\text{oper}}[k] > \epsilon_{\text{oper}})}$ **then**
        `// New data adds information and is operationally relevant`
        return Accept $\mathbf{c}_k$
    else
        `// New data does not improve informativity or is not operationally relevant`
        return Reject $\mathbf{c}_k$
    end
end

The three indicators address distinct aspects of data quality:

The interplay between the three indicators enables diagnosis of various operational conditions. The diagnostic capability proves particularly valuable during seasonal transitions in building control, where outdoor conditions and occupancy patterns shift significantly. For instance, the model degradation scenario frequently occurs during the first cold days of autumn, when the system encounters temperature ranges not present in summer training data.

3 Adaptive Strategy Validation

3.1 Numerical Case Study: RC Parameter Variance

This numerical case study is an open-loop validation. This numerical case study uses a simplified thermal system to demonstrate adaptive prediction mechanisms. A second-order RC network provides clear observation of data matrix evolution and update strategy behaviors. This controlled environment removes confounding effects from nonlinearities and high-frequency disturbances. The reduced

Table IV.1: Interpretation of data quality indicators: high vs low values

Case	Explanation
Model Matching Indicator $I_{\text{match}}[k]$	
High ($\geq \bar{I}_{\text{match}}$)	Current model provides good predictions. System behavior remains consistent with historical data. Data matrix represents current dynamics adequately.
Low ($< \bar{I}_{\text{match}}$)	Prediction errors exceed expected levels. System dynamics have drifted from historical behavior. Model requires refresh with new measurements to recover accuracy.
Data Informativity Indicator $I^{\text{aug}}_{\text{info}}[k]$	
High ($\geq \epsilon_{\text{info}}$)	Input trajectories provide persistent excitation. Data matrix has full numerical rank for reliable prediction. New data expands subspace representation.
Low ($< \epsilon_{\text{info}}$)	Input trajectories lack sufficient diversity. Data matrix has rank deficiency. New measurements do not excite system directions needed for robust control.
Operational Relevance Indicator $I_{\text{oper}}[k]$	
High ($I_{\text{oper}}[k] = 1$)	Data collected during service hours when control performance is critical. Represents typical operating conditions. Directly improves model for intended use.
Low ($I_{\text{oper}}[k] = \alpha_{\text{off}} < 1$)	Data collected during off-service periods (nights, weekends). Limited relevance for service-hour control performance. May introduce spurious patterns unrelated to normal operation.

computational burden enables exploration of forgetting factors, window sizes, and rank conditions. These findings inform subsequent controller design for realistic building benchmarks.

Components in algorithm are examined step by step in this case study, which is organized as follows:

1. **System formulation:** The numerical test case formulation is introduced with a parameter shift at a known time. Dataset perfectness is confirmed by demonstrating that Willems' lemma predictor achieves perfect prediction under ideal conditions before the parameter change occurs.

2. **Sliding window vs. recursive reduction in open-loop validation:** Sliding data window and recursive reduction strategies are applied to open-loop validation without data quality indicators. All new data is accepted. The results reveal different trade-offs: sliding window exhibits sharp transient peak errors during parameter transitions but achieves zero error prediction, while recursive reduction maintains smoother error trajectories with residual errors in long term.

3. **Forgetting factor impact on the peak-residue trade off:** A forgetting factor is incorporated into the recursive reduction method. The results confirm that the trade-off between transient peak error and residue error depends critically on the effective model memory length, which is controlled by the forgetting factor γ. Sensitivity analysis quantifies how γ selection influences adaptation dynamics.

4. **Persistent excitation degradation and quality assessment:** All three update strategies are evaluated with constant input. This experiment demonstrates the critical importance of data quality indicators in preventing matrix rank deficiency and control performance degradation when new measurements lack sufficient information content.

The insights from this controlled numerical analysis provide concrete guidance for configuring the three adaptive mechanisms in the MOSA benchmark validation presented in Section 3.2.

3.1.1 Second-order Thermal System with Parameter Shift

The case study uses a second-order RC thermal network representing two thermally coupled zones, as shown in figure IV.1. As established in chapter II, this configuration captures heat transfer dynamics through resistive-capacitive elements. Each zone exchanges heat with the external environment through thermal resistance $R_{e,i}$ and stores thermal energy in capacitance C_i. The inter-zone coupling is characterized by thermal resistance $R_{1,2}$. Heating power inputs P_1 and P_2 serve as control variables and zone temperatures T_1 and T_2 constitute the system outputs.

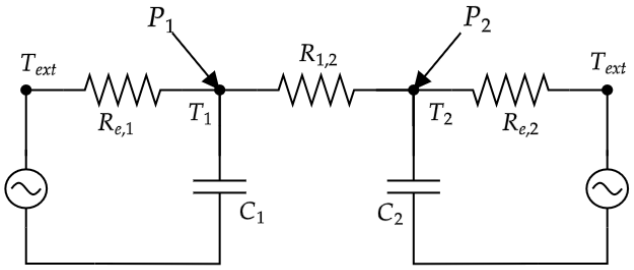

Figure IV.1: Second-order RC thermal network topology. Two thermally coupled zones with external resistances $R_{e,1}$ and $R_{e,2}$, inter-zone resistance $R_{1,2}$. Heating powers P_1 and P_2 are control inputs, zone temperatures T_1 and T_2 are system outputs.

The system dynamics follow the ordinary differential equations:

$$\frac{dT_1}{dt} = \frac{1}{R_{e,1}C_1}(T_{ext} - T_1) + \frac{1}{R_{1,2}C_1}(T_2 - T_1) + \frac{1}{C_1}P_1 \tag{123a}$$

$$\frac{dT_2}{dt} = \frac{1}{R_{e,2}C_2}(T_{ext} - T_2) + \frac{1}{R_{1,2}C_2}(T_1 - T_2) + \frac{1}{C_2}P_2 \tag{123b}$$

The input vector is $\mathbf{u} = [P_1, P_2]^\top$ and the output vector is $\mathbf{y} = [T_1, T_2]^\top$. The external temperature T_{ext} is held constant to isolate internal dynamics.

The continuous-time model is discretized using fourth-order Runge-Kutta (RK4) integration with sampling period $\Delta t = 600$ seconds. This time step provides adequate temporal resolution while

maintaining numerical stability. The discretization preserves system linearity and enables direct application of Willems' fundamental lemma.

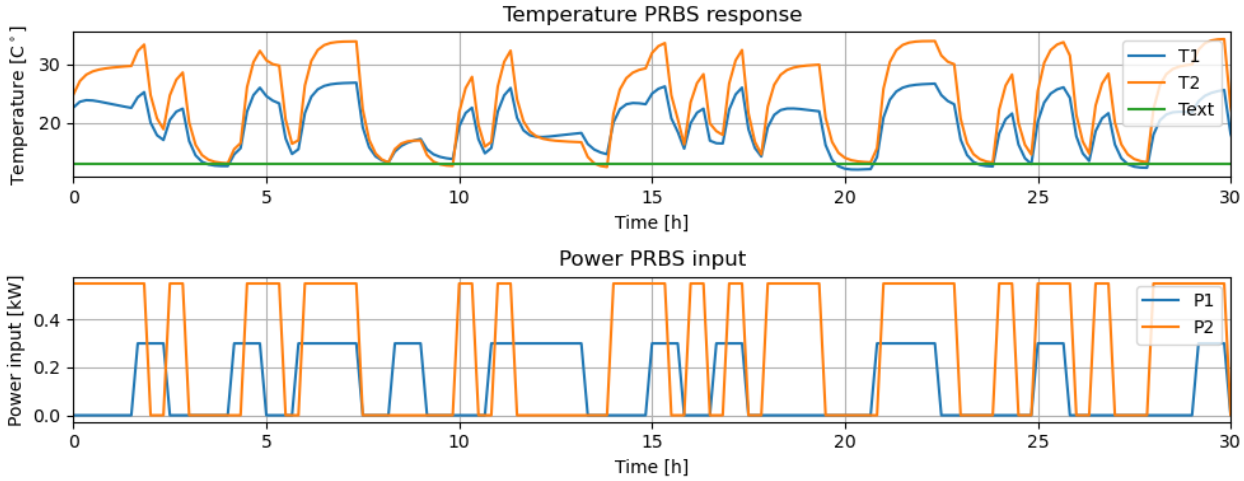

Figure IV.2: PRBS excitation dataset for initial Hankel matrix construction. Top panels: zone temperatures T_1 and T_2. Bottom panels: heating power inputs P_1 and P_2. The 30-hour dataset ensures full rank condition for data-driven prediction.

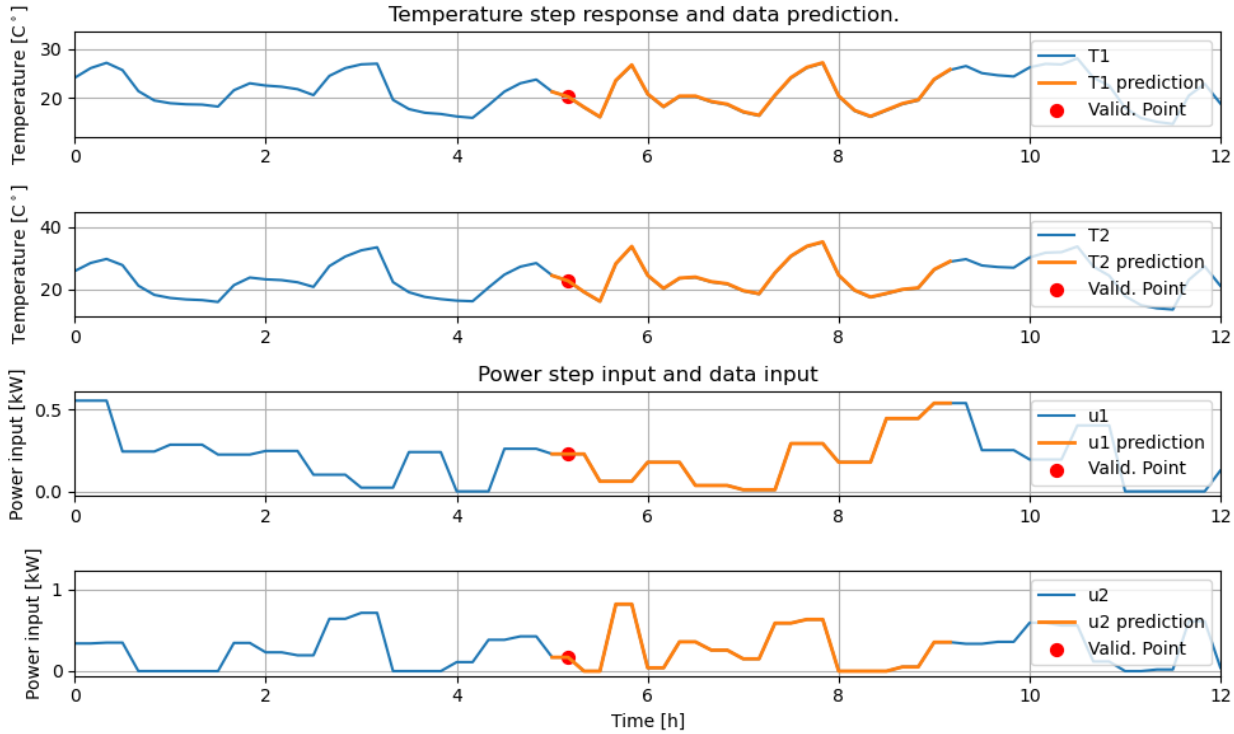

Figure IV.3: Validation of data-driven prediction under constant parameters. Blue curves: actual system trajectories. Orange curves: predictions from validation points (red dots). The model achieves perfect accuracy with $T_{ini} = 2$ and $H = 24$ steps.

A dataset spanning 30 hours is generated using pseudo-random binary sequence (PRBS) excitation applied to both heating inputs, as shown in figure IV.2. The PRBS signal ensures persistent excitation across the system's frequency range, satisfying the rank condition for data-driven prediction. To evaluate adaptive capabilities, the external thermal resistances $R_{e,1}$ and $R_{e,2}$ are doubled at $t = 32$ hours, introducing an abrupt parameter shift. This discontinuity creates distinct "old system" and "new system" regimes.

The validation experiment uses a randomized input signal distinct from the training PRBS sequence. Figure IV.3 shows the data-driven model performance with initialization horizon $T_{ini} = 2$ steps and prediction horizon $H = 24$ steps (6 hours). The choice $T_{ini} = 2$ corresponds to the second-order dynamics of the continuous-time system in equations (123). Red dots mark validation points where predictions are initiated, with orange trajectories showing the predicted evolution.

The objective is to confirm that the initial PRBS dataset provides ideal prediction capability before introducing parameter variations. Perfect alignment between predictions and actual responses verifies that Willems' fundamental lemma conditions are satisfied for this linear time-invariant system under persistent excitation.

3.1.2 Comparison of Performance Under System Transition

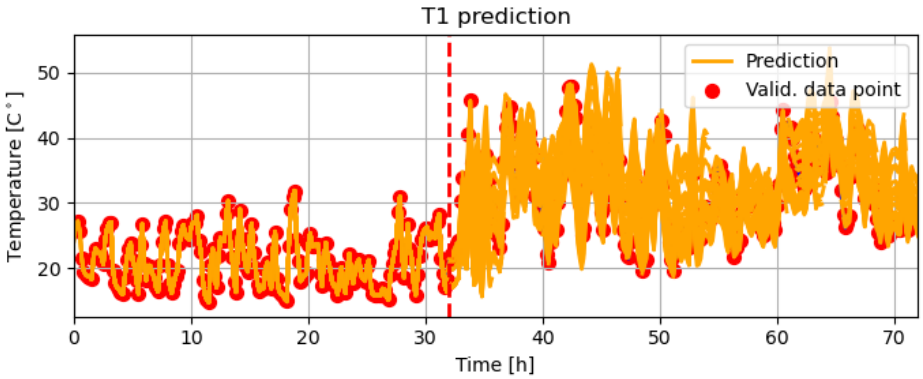

Figure IV.4: Non-adaptive DeePC prediction under abrupt parameter shift at $t = 32$ hours (marked by red dashed line). Zone 1 is shown as an example, the two rooms' results is comparable. The prediction failed at the shifting moment. Red scatters are actual zone temperature sample in dataset. Yellow trajectories: predictions from static dataset. Prediction accuracy degrades significantly after parameter change.

Figure IV.4 shows the prediction behavior of static DeePC when the system undergoes the parameter shift. Before the shift, predictions align perfectly with actual trajectories, confirming the validity of the initial data matrix. After $t = 32$ hours, prediction accuracy degrades progressively as the horizon extends.

Figure IV.5 quantifies this degradation through mean squared error evolution. The MSE remains near zero during the first 32 hours when system parameters match the training dataset. Following the parameter shift, prediction errors increase substantially and show persistent oscillations without convergence. This demonstrates that outdated data matrices cannot maintain prediction accuracy under parameter variations.

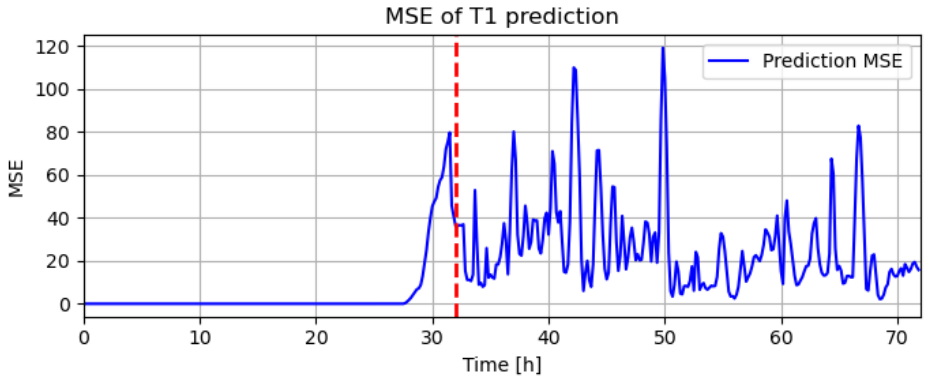

Figure IV.5: Mean squared error evolution for non-adaptive prediction. Parameter shifts on $t = 32$ hours (red dashed line).

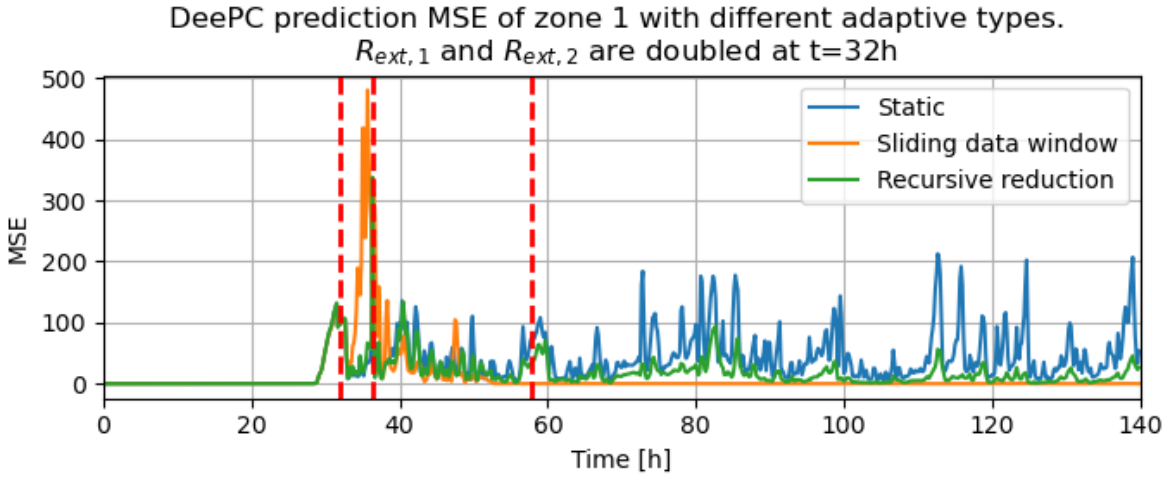

Figure IV.6: Mean squared error comparison for adaptive update strategies (mechanism configuration see table IV.2). Red dashed lines: parameter shift ($t = 32$ h), peak error ($t \approx 36$ h), and complete window renewal ($t \approx 58$ h). All new column is used without filtered according to indicators.

Table IV.2: Adaptive mechanism embedded in figure IV.6 experiments

Curve	**Sliding data window**	**Recursive reduction**	**Forgetting factor**	**Indicators**
Blue	–	–	–	–
Orange	✓	–	–	–
Green	–	✓	–	–

Figure IV.6 compares prediction accuracy across three strategies. The static approach maintains the original 30-hour dataset without updates. The sliding window strategy continuously removes the oldest data column and appends new measurements, maintaining fixed Hankel dimensions. The recursive reduction approach incrementally incorporates new data through fast SVD updates, compressing the growing column space.

The sliding window approach shows a significant error peak immediately following the parameter shift at $t = 32$ hours, reaching MSE values exceeding 100 for both zones. This peak arises because newly appended columns contain transient dynamics from the parameter shift itself, temporarily degrading the data matrix quality. After approximately 25 hours, the error converges toward zero as outdated pre-shift data is completely purged from the Hankel matrix. This convergence timeline corresponds to the window length.

The recursive reduction method maintains substantially lower prediction errors throughout the transition period, avoiding the dramatic peak observed in the sliding window approach. While the MSE increases moderately after $t = 32$ hours and remains persistently above zero, the error magnitude stays well below the static baseline. The absence of sharp peaks indicates that the SVD-based compression dilutes the influence of transient shift dynamics across the reduced basis. However, without explicit forgetting, historical pre-shift information persists indefinitely in the compressed representation, preventing complete adaptation to the new parameter regime.

The temporal evolution can be segmented into four distinct periods:

1. **Pre-shift Period ($t = 0$ to $t = 32$ hours):** All three strategies maintain excellent prediction accuracy with MSE near zero throughout the initial period. The linear system operates under constant parameters, and the persistent excitation in the training dataset ensures full rank conditions. Prediction errors begin rising from $t = 28$ hours onward, which is an expected behavior. The prediction horizon is 4 hours, so at $t = 28$ hours, the prediction validation moving window begins to include system behavior after the parameter shift time at $t = 32$ hours. Since the parameter shift is not informed to model, the rising errors during the period from $t = 28$ to $t = 32$ hours are therefore inevitable.

2. **Critical transition period ($t = 32$ to $t \approx 36.3$ hours):** Differences between update strategies emerge. The sliding window approach shows a pronounced error peak as newly appended data columns capture the parameter shift transient itself. The peak reaches maximum magnitude at approximately $t = 32 + (T_{ini} + H)\Delta t$. Beyond this point, columns entering the Hankel matrix characterize purely the new system regime. The recursive reduction strategy avoids this peak because the SVD compression distributes transient effects across the reduced basis, preventing any single corrupted column from dominating the prediction.

3. **Asymptotic adaptation period ($t \approx 36.3$ to $t = 57.8$ hours):** The two adaptive strategies converge to distinct steady-state behaviors. The sliding window approach achieves zero prediction error at approximately $t \approx 57.8$ hours, calculated as $t = 32 + n_{data} \cdot \Delta t$, where n_{data} represents the total number of columns in the Hankel matrix. At this instant, all pre-shift data has been completely purged from the dataset. The recursive reduction method shows gradually decreasing errors as the proportion of new system data grows, diluting the residual influence of outdated information. However, without an explicit forgetting mechanism, this strategy cannot eliminate historical data entirely, resulting in persistent low-level prediction errors.

The contrasting behaviors stem from their different approaches to information retention. The recursive reduction method constructs a low-dimensional subspace approximation through singular value decomposition, projecting the Hankel matrix onto its principal components. During the incremental update process, all existing column space information contributes equally, regardless of temporal recency. Consequently, the predictor incorporates behavioral characteristics from both

pre-shift and post-shift regimes. In contrast, the sliding window approach implements strict temporal bounds by explicitly removing the oldest columns from the Hankel matrix. This hard truncation mechanism completely eliminates outdated column space information, enabling the predictor to converge toward perfect accuracy once all pre-shift data has been purged.

The distinction lies in memory management: reduction preserves a weighted combination of all historical information within the compressed representation, while windowing enforces finite memory through deletion. The sliding window strategy achieves complete adaptation by purging outdated information, but suffers pronounced error peaks during transition periods. The recursive reduction approach maintains smoother error trajectories by distributing new information across a compressed basis, yet cannot fully eliminate historical influence without explicit forgetting.

3.1.3 Forgetting Mechanism for Enhanced Adaptivity

To verify that memory management determines adaptation behavior, an intermediate strategy is introduced that combines characteristics of both approaches. This method modifies the recursive SVD algorithm by exponentially weighting historical information through multiplicative scaling of singular values. Before incorporating a new data column at time step $k+1$, the existing singular value matrix is scaled by forgetting factor γ in equation (107). The forgetting factor determines effective memory length through $N_{eff} = \lceil \ln \epsilon_w / \ln \gamma \rceil$.

This represents an intermediate approach because it privileges recent measurements while maintaining continuous representation of historical data. Past column space contributions fade exponentially but never reach exactly zero, avoiding the hard truncation of sliding windows.

Figure IV.7 compares four forgetting factors: $\gamma = 1.0$ (no weighting), and $\gamma = 0.999, 0.99, 0.98$. Decreasing γ creates a trade-off between transient and steady-state errors. Recursive reduction with forgetting achieves near-perfect steady-state accuracy, confirming that exponential weighting removes outdated system characterization.

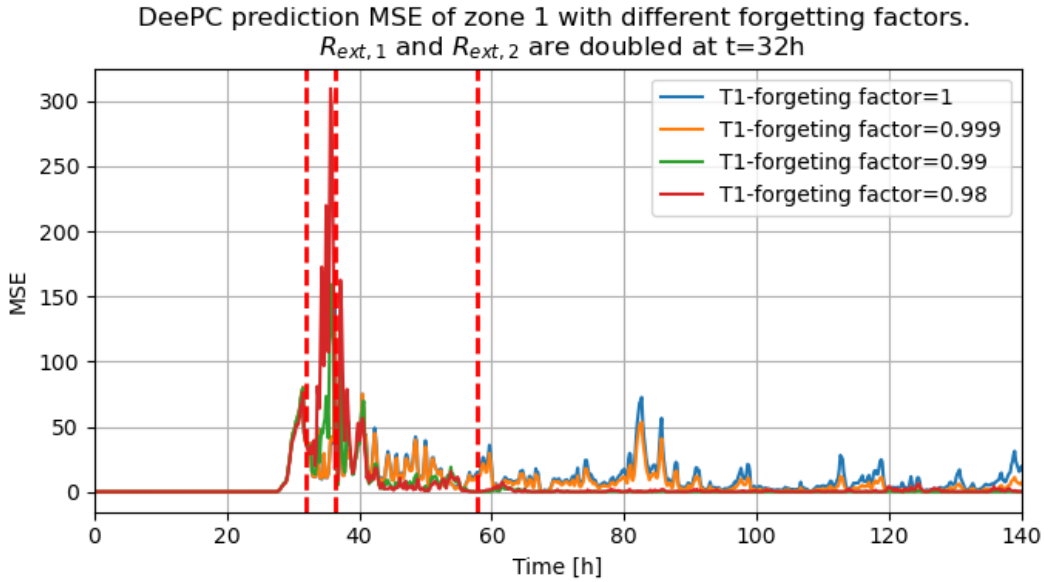

Figure IV.7: Sensitivity of adaptation behavior to forgetting factor selection (mechanism configuration see table IV.3).

Table IV.3: Adaptive mechanism embedded in figure IV.7 experiments

Curve	Sliding data window	Recursive reduction	Forgetting factor	Indicators
Blue	–	✓	– ($\gamma = 1.0$)	–
Orange	–	✓	Weak($\gamma = 0.999$)	–
Green	–	✓	Moderate($\gamma = 0.99$)	–
Red	–	✓	Aggressive($\gamma = 0.98$)	–

The convergence mechanism is quantified through accumulated weighting. The Hankel matrix contains 155 columns. After 155 updates with $\gamma = 0.99$, original information is attenuated by $\gamma^{155} \approx 0.21$ (21% remains). This decay explains why forgetting reduction approaches zero error: pre-shift information becomes negligibly small relative to post-shift measurements.

Reducing γ amplifies the transient peak during parameter shift while improving asymptotic convergence. The case $\gamma = 0.98$ shows the largest peak but fastest error decay. The case $\gamma = 0.999$ exhibits smoother transition with slower decay. The intermediate choice $\gamma = 0.99$ balances both effects, producing manageable transient peaks with reliable convergence.

Recursive reduction with $\gamma = 0.99$ reproduces the error pattern of sliding window approaches. This confirms that forgetting—whether through column deletion or singular value weighting—dominates peak magnitude and steady-state behavior. Sliding window enforces finite memory through column purging. Recursive reduction provides greater flexibility: γ independently tunes memory length while maintaining constant matrix size. The SVD procedure also incorporates implicit quality assessment through rank verification.

3.1.4 Persistent Excitation Degradation Scenario

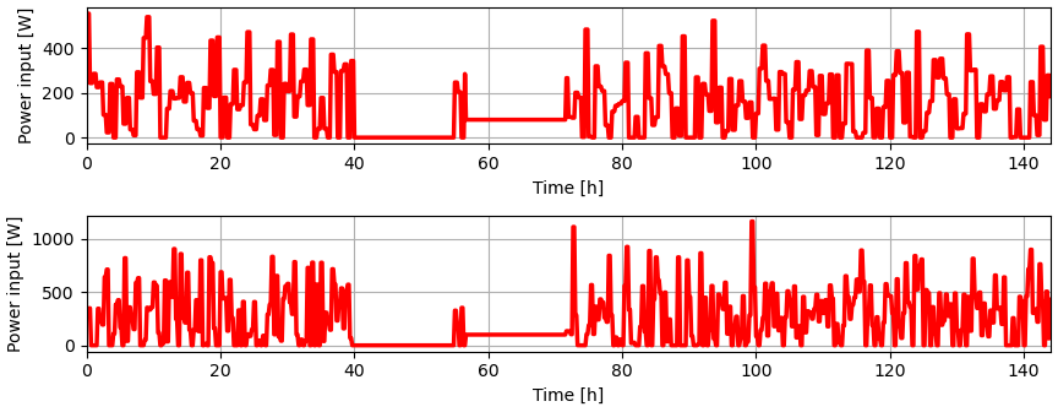

Figure IV.8: Validation input signals with persistent excitation degradation periods. The input signal is still based on the randomized signal but replace two period by constant input. Top and bottom panels: heating power inputs P_1 and P_2. Two constant input stages (shaded regions): $t = 40$-55 hours and $t = 56$-64 hours.

The preceding experiments used continuously varying input signals that maintained persistent excitation throughout the validation period. However, due to input saturation, closed-loop building

control frequently shows quasi-constant actuator commands over extended periods as the system approaches thermal equilibrium or maintains setpoint conditions. To evaluate adaptive performance under degraded excitation, this experiment incorporates two deliberate constant input stages, as shown in figure IV.8. The first constant phase spans from $t = 40$ to $t = 55$ hours, followed immediately by a second stage from $t = 56$ to $t = 72$ hours. Both periods are positioned within the asymptotic adaptation phase, where residual information from the pre-shift system has not yet been completely purged from the data matrix.

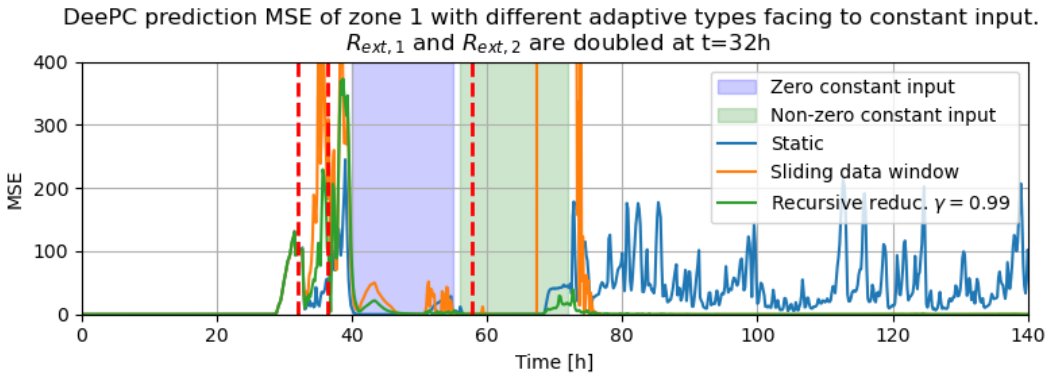

Figure IV.9: Mean squared error evolution under persistent excitation degradation for baseline adaptive strategies (mechanism configuration see table IV.4). Red dashed lines: parameter shift ($t = 32$ h), first constant input stage ($t = 40$-55 h), second constant input stage ($t = 56$-72 h). Red dashed line marks the critical moment as before.

Table IV.4: Adaptive mechanism embedded in figure IV.9 experiments

Curve	Sliding data window	Recursive reduction	Forgetting factor	Indicators
Blue	–	–	–	–
Orange	✓	–	–	–
Green	–	✓	Moderate($\gamma = 0.99$)	–

Figure IV.9 reveals differences in how each strategy responds to persistent excitation degradation. The sliding data window approach suffers significant performance collapse when the Hankel matrix becomes populated with constant input measurements. During the first constant stage ($t = 40$-55 h), the rank condition deteriorates progressively as non-informative columns replace well-excited historical data. The impact shows most severely after $t = 64$ hours in sliding data window case, where prediction errors exceed MSE = 400 for both zones. In contrast, recursive reduction strategies maintain substantially lower error magnitudes throughout both constant input periods. The fast-SVD update process inherently evaluates singular values when incorporating new columns, effectively attenuating the contribution of poorly-excited measurements in the compressed representation.

To address persistent excitation degradation, the data quality indicators introduced in section 2.2 are integrated into both sliding window and recursive reduction strategies. For this case study validation, the operational relevance indicator I_{oper} is deactivated since no occupancy patterns are involved. The column acceptance logic combines both indicators through the decision rule:

$$\text{Accept new column} \iff (I_{\text{info}}^{\text{aug}} > 0.97) \vee (I_{\text{match}} > 0.3) \tag{124}$$

where $I_{\text{info}}^{\text{aug}}$ denotes the informativity of the augmented Hankel matrix and I_{match} quantifies recent prediction error magnitude. This formulation accepts new measurements either when they improve matrix rank condition or when prediction accuracy has degraded sufficiently to warrant dataset refresh.

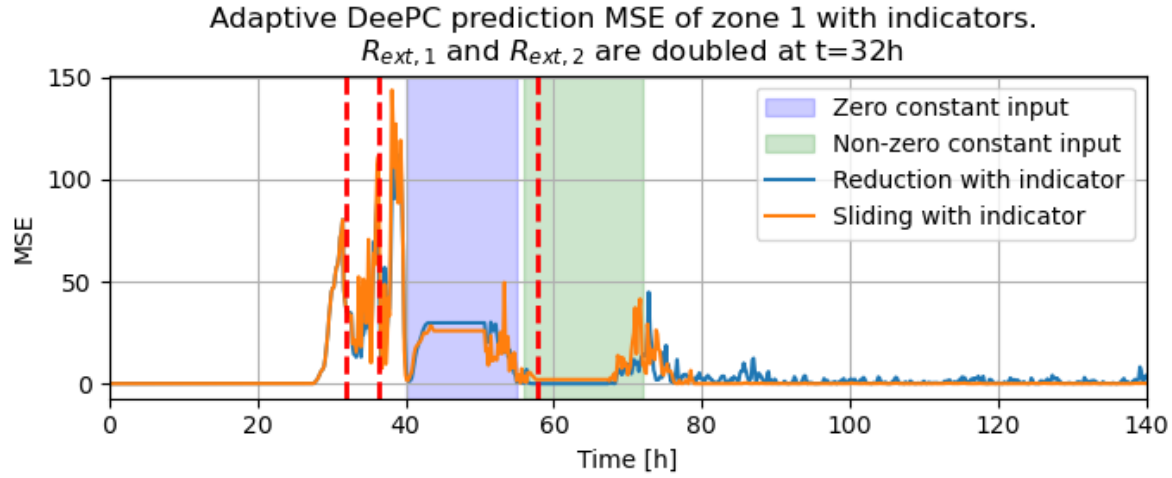

Figure IV.10: Performance comparison with data quality indicator integration under persistent excitation degradation (mechanism configuration see table IV.5). Red dashed lines: parameter shift ($t = 32$ h) and constant input periods ($t = 40$-55 h, $t = 56$-64 h).

Table IV.5: Adaptive mechanism embedded in figure IV.9 experiments

Curve	Sliding data window	Recursive reduction	Forgetting factor	Indicators
Blue	✓	–	–	✓
Orange	–	✓	Moderate($\gamma = 0.99$)	✓

Figure IV.10 shows that incorporating these quality checks transforms the behavior of both adaptive strategies. The significant error peaks observed in baseline implementations are eliminated, and asymptotic convergence toward low steady-state errors is achieved. The mechanism is straightforward: informativity verification actively rejects constant input measurements during both degraded excitation periods, preventing the accumulation of rank-deficient columns in the Hankel matrix. Compared to the failure of sliding data window, the indicator-enhanced approach maintains errors below MSE = 20 throughout, demonstrating substantial robustness improvement under challenging excitation conditions.

3.1.5 Summary

The case study validation examined adaptive data-driven prediction under controlled conditions. Temporal segmentation into four distinct phases enabled precise characterization of how sliding window and recursive reduction strategies respond to abrupt parameter shifts. The forgetting factor

investigation quantified this compromise, establishing that exponential weighting with $\gamma = 0.99$ achieves balanced performance by combining manageable transition peaks with reliable asymptotic convergence.

The persistent excitation degradation experiment extended validation beyond idealized conditions to evaluate robustness under informationally deficient measurements. Baseline sliding window implementations suffered significant failure (MSE $>$ 400) as non-informative columns displaced well-excited historical data. Integration of data quality indicators transformed this behavior by actively rejecting rank-deficient measurements. The indicator-enhanced strategies maintained prediction errors below MSE = 20 throughout both constant input periods.

The simplified RC thermal network provides an idealized testbed for isolating adaptation mechanisms. However, realistic building thermal systems show continuous parameter variations driven by seasonal climate shifts, gradual equipment degradation, and evolving occupancy patterns rather than abrupt discontinuities. To validate the generalizability of these findings, the subsequent subsection evaluates adaptive strategies on the MOSA benchmark.

3.2 MOSA benchmark validation

The MOSA benchmark validation maintains identical building geometry, HVAC system specifications, and baseline controller configurations established in previous chapters. The five-zone office building model operates under the same weather forcing, occupancy schedules, and thermal comfort constraints. The adaptive DeePC controller extends the static configuration from Chapter 3 by integrating the proposed adaptive algorithm (see algorithm 1). Table IV.6 summarizes the adaptive-specific parameters, while all other controller settings remain unchanged from the baseline DeePC implementation.

Table IV.6: Adaptive DeePC controllers configuration for MOSA benchmark. Adaptive strategy is algorithm 1. DeePC and simulation configuration follows table III.1

Parameter	**Symbol**	**Value**
Forgetting factor	γ	0.99
Prediction error threshold	$\bar{I}_{\text{pred}}$	0.3
Informativity threshold	ϵ_{info}	0.985
Rank tolerance (relative)	ϵ_{tol}	10^{-2}
Off-service weight	α_{off}	0.1
Service hours	$\mathbb{T}_{\text{service}}$	10:00–16:00

Figure IV.11 demonstrates closed-loop prediction quality under realistic operational conditions spanning three days of MOSA benchmark operation. The top panel presents zone temperature evolution. Blue curves indicate measured thermal states and overlaid gray trajectories show multi-step predictions initiated at regular intervals throughout the control horizon. The predicted trajectories consistently converge toward actual measurements despite continuously varying weather disturbances and control actions, confirming effective online adaptation to evolving system dynamics.

The middle panel displays thermal power consumption, showing characteristic diurnal patterns with morning heating demands and afternoon passive operation under solar gains. The bottom panel

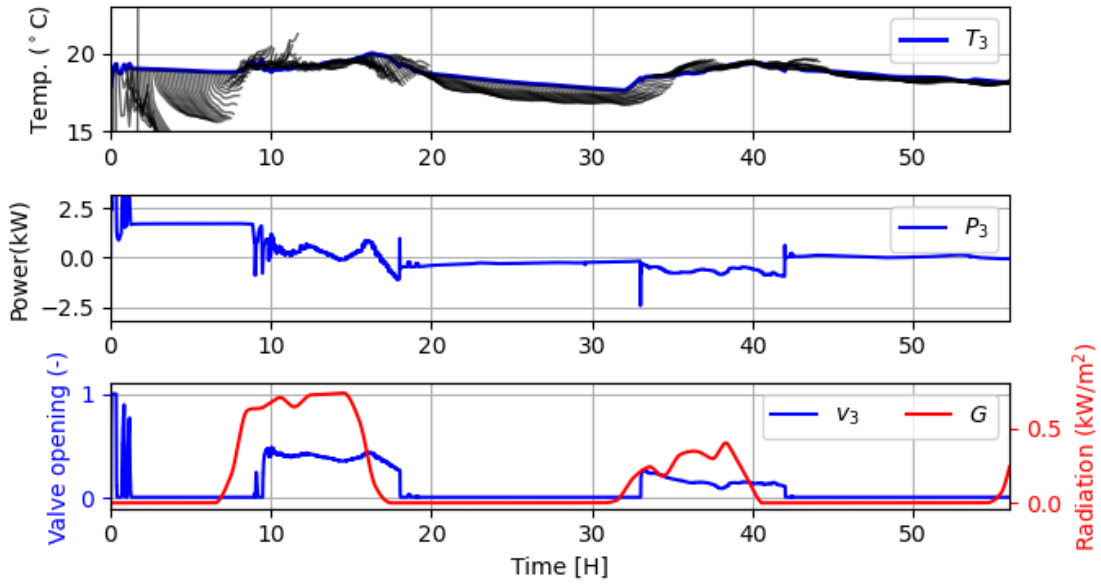

Figure IV.11: Closed-loop prediction quality validation over three-day MOSA benchmark operation. The north zone result is shown in figure, other zones' behavior are similar. Top: zone temperature trajectories with predictions (gray) and measurements (blue). Middle: thermal power consumption showing diurnal patterns. Bottom: valve actuation (blue) and solar radiation disturbance (red).

tracks direct system inputs—valve opening commands (blue) and incident solar radiation (red)—which serve as measured disturbances to the predictive model.

Extended periods of quasi-constant valve operation occur during mid-morning ($t = 10$-20 h) and afternoon ($t = 30$-45 h) intervals when the system maintains near-equilibrium conditions. Despite these persistent excitation deficiencies that would degrade baseline implementations, the prediction trajectories continue converging accurately to measured states. This validates that the integrated quality indicators successfully prevent matrix degradation by rejecting informationally deficient measurements while preserving well-excited historical data sufficient for reliable prediction.

Figure IV.12 presents Pareto front comparisons across different weighting configurations, following the multi-objective evaluation methodology established in previous chapters. Both sliding window and recursive reduction achieve clearly better trade-offs than static DeePC, reducing energy consumption across the full range of comfort violation levels.

The two adaptive approaches show comparable overall performance. The sliding data window with column quality indicators achieves a slightly better balance between energy consumption and comfort, producing a more favorable front in the low-to-moderate violation range. This advantage is attributed to the selective data incorporation mechanism, which filters out low-quality measurements and preserves excitation conditions in the Hankel matrix. These results confirm that online data adaptation is the important factor of performance improvement. The sliding window approach offering a practical advantage when combined with quality-based data filtering.

It should be noted that these results are specific to the MOSA benchmark system. For other building types or configurations, the relative performance of each strategy may differ due to variations in thermal dynamics, occupancy patterns, and operation principle. However, the component-level analysis and hyper-parameter studies presented in this chapter provide a structured methodology to

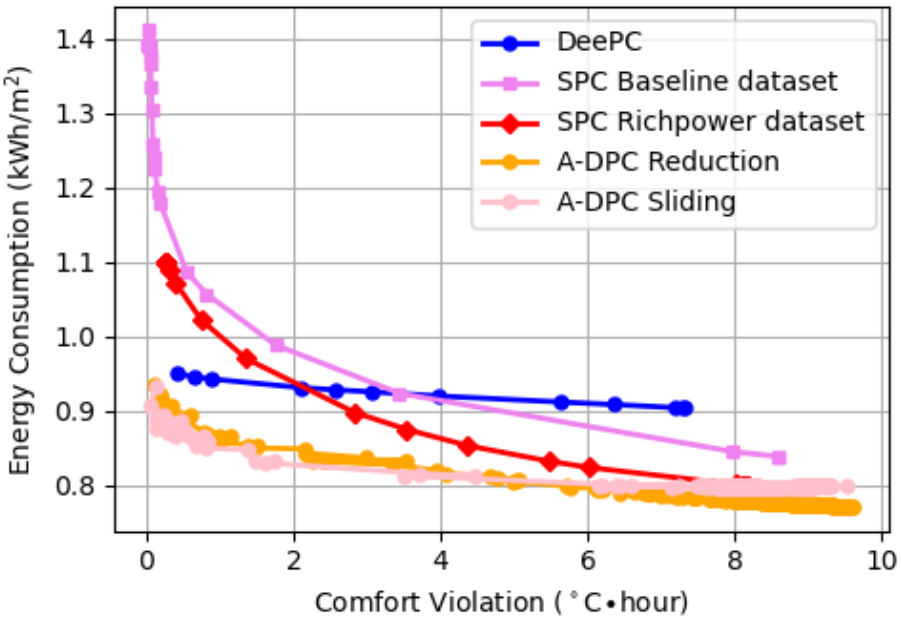

Figure IV.12: Pareto front comparison of adaptive and baseline predictive control strategies on MOSA benchmark. Blue: static DeePC. Pink: SPC with baseline dataset. Red: SPC with rich-power dataset. Orange: Adaptive DeePC with reduction configuration in table IV.4. Light pink: Adaptive DeePC with sliding data window configuration in table IV.4.

evaluate each design choice in a given application context. This serves as a practical guide for both engineers deploying adaptive predictive control and researchers investigating new update strategies.

4 Chapter Conclusion and Outlook

This chapter developed an adaptive data-driven predictive control framework for building thermal management. The central motivation is that building systems exhibit time-varying dynamics due to seasonal changes, occupancy shifts, and equipment aging, which cause static data matrices to lose prediction accuracy over time. The chapter addressed this challenge through online Hankel matrix update strategies, forgetting mechanisms, and data quality indicators.

Two update paradigms were studied: sliding window methods, which maintain fixed matrix dimensions by replacing old columns with new measurements, and recursive reduction approaches, which compress a growing dataset using fast SVD updates. Both strategies effectively track parameter changes, with complementary characteristics. Sliding window approaches achieve asymptotic convergence after a full window renewal, while recursive reduction maintains smoother error trajectories during transitions but retains residual influence from outdated data. Exponential temporal weighting with $\gamma = 0.99$ improves both strategies by gradually reducing the influence of older measurements without introducing hard truncation.

While individual metrics for prediction accuracy and rank analysis exist in the literature [156, 157], their integration into a unified framework with a third indicator for operational relevance represents an original approach. Three indicators address distinct aspects of measurement utility: prediction error

(I_{match}) detects accuracy degradation, informativity (I_{info}) verifies persistent excitation through rank analysis, and operational relevance (I_{oper}) prioritizes measurements from occupied periods. When baseline strategies blindly incorporate rank-deficient measurements, significant matrix degradation occurs. The indicator-enhanced strategies prevent this failure by selectively rejecting uninformative columns, maintaining errors below MSE $= 20$ under the same conditions.

Validation on the RC thermal network isolated these mechanisms under controlled conditions. Validation on the MOSA benchmark confirmed that these improvements generalize to realistic scenarios. Both adaptive DeePC variants achieve clearly better Pareto fronts than static DeePC, reducing energy consumption across all comfort violation levels. The two adaptive strategies show comparable overall performance, with the sliding window approach combined with quality indicators achieving a slightly better energy-comfort trade-off.

The adaptive framework confirms that online data management is the key driver of performance improvement over static baselines. The component-level analysis and hyper-parameter studies in this chapter provide a structured methodology applicable to different building systems.

Future directions include extending the adaptation framework to nonlinear dynamics through kernel-based methods, developing theoretical guarantees for closed-loop stability under online updates, and investigating strategies for automatic threshold selection for the quality indicators across different building types.

Chapter V

Thesis Conclusions and Perspectives

1 Synthesis and Contributions 132
1.1 Hierarchical Framework and Validation Protocol 132
1.2 Data-Driven Predictive Control Approaches 133
1.3 Adaptive Mechanisms 135
1.4 Thesis Conclusion 136
1.5 Valorisation and Scientific Contributions 136
2 Perspectives 137
2.1 Perspectives on Data-Driven Predictive Control 137
2.2 Perspectives on Building Control Implementation 138

1 Synthesis and Contributions

This thesis explores data-driven predictive control approaches for building thermal management. Building energy consumption is critical for decarbonization objectives, but advanced control faces practical barriers. Two main challenges exist: equipment diversity across different HVAC systems, and limited training data from normal building operation. This research investigates how data-driven control methodes address these challenges through comparative analysis and standardized validation.

This thesis proposes a hierarchical control framework for building thermal management. The framework separates control into two levels: lower-level controllers track heat power setpoints, while an upper-level controller optimizes thermal setpoints. The work compares three different predictive control methods: physics-based models (RC-MPC), subspace identification (SPC), and direct data-based control (DeePC). These methods show different trade-offs: they require different amounts of data, offer different control quality, and demand different computational resources. Adaptive versions of these methods are also developed to handle changing weather and equipment performance over time.

1.1 Hierarchical Framework and Validation Protocol

The thesis proposes a hierarchical control framework that addresses equipment diversity and data limitations in building thermal management. The architecture separates control into two levels: the supervisory level manages long-term thermal planning and optimization; the tracking level executes short-term equipment regulation. This separation provides equipment-agnostic supervisory control that operates on thermal power abstractions rather than actuator positions, enabling deployment across diverse HVAC configurations.

At the supervisory level, Model Predictive Control (MPC) is used to optimize thermal comfort and energy consumption over a prediction horizon of several hours. MPC exploits the building's thermal inertia by planning power setpoints in advance, enabling better energy management. The supervisory controller uses a long sampling time (typically 15–30 minutes) to reduce computational burden while remaining responsive to slow thermal dynamics.

At the tracking level, Virtual Reference Feedback Tuning (VRFT) is employed to tune power-following controllers without requiring explicit HVAC system models. This is the first application of VRFT to building thermal management. VRFT optimizes controller parameters directly from operational data, making it applicable to diverse equipment types. The tracking controllers use a fast sampling time (a few minutes) to follow power setpoints accurately while compensating for equipment dynamics.

Validation uses the BOPTEST standardized simulation benchmark. Two test cases were selected to highlight the universality of the proposed layered control methodology, demonstrating its applicability regardless of the specific equipment under consideration. MOSA features air handling systems (typical in large buildings), while MZRH features hydronic heating (typical in residential buildings). These test cases cover the main equipment types encountered in building thermal management. Performance is assessed through Pareto frontier analysis, quantifying the energy–comfort trade-offs achieved by each control approach.

1.2 Data-Driven Predictive Control Approaches

The supervisory level investigation compares indirect control via system identification against direct control without explicit models.

The indirect approach compares physics-informed gray-box RC models against data-driven black-box subspace identification under limited data. Validation reveals comparable open-loop prediction accuracy but divergent closed-loop performance. The gray-box approach demonstrates superior extrapolation through energy conservation structure, achieving substantial energy savings while maintaining comfort. Physical structure constrains predictions to thermodynamically feasible trajectories, enabling reliable operation beyond training data. The rate-constrained formulation explicitly models gradual thermal power evolution, improving energy management without compromising comfort.

The direct approach investigates whether model-free predictive control can eliminate explicit system identification while maintaining performance. DeePC constructs predictions directly from input-output trajectories through Hankel matrix operations, bypassing model structure selection entirely. Indirect approaches extract fixed models offline through subspace projection, while DeePC determines models online through regularized trajectory combination.

Benchmarking demonstrates DeePC achieves superior closed-loop performance, reducing energy consumption and thermal discomfort while eliminating modeling procedures. Hyperparameter analysis establishes practical guidelines for regression order selection, regularization configuration, and computational considerations. Table V.1 summarizes the evaluation of the three predictive control approaches.

RC-MPC relies on a physics-informed model structure where energy conservation is enforced by the RC network topology. This structural constraint ensures inherent model stability and offers robust closed-loop performance with minimal tuning and low computational demands. As shown in the Chapter 2 case study, this property gives RC-MPC a clear advantage over SPC when the training dataset

Table V.1: Comparative evaluation of data-driven predictive control approaches for building thermal management.

Criterion	**RC-MPC**	**SPC**	**DeePC**
Modeling Effort	High	Low	Low
Data Requirement	Medium	Medium	High
Development Complexity	Medium	Medium	High
Online Optimization Time	Small	Medium	Medium
Computation Memory	Small	Small	High
Tuning Effort	Low	Medium	High
Achievable Performance	High	Medium–High	High

is collected under closed-loop operation with limited input excitation. The RC model identification is performed via least-squares estimation, which is computationally efficient, closed-form optimal, and scales well with the number of zones. The model structure can also be configured according to available building knowledge (like zone segmentation, thermal order per zone, and HVAC interface). Finally, the convexity of both the identification and the control optimization problems keeps the tuning effort low compared to the other approaches in this study, making RC-MPC a practical and reliable baseline for building thermal management. RC-MPC is recommended when physical modeling expertise and domain knowledge are available.

SPC is a convenient approach for implementation, as subspace identification methods such as N4SID are well-established and widely used across engineering domains. However, subspace identification is more sensitive to dataset quality than RC modeling: when data is collected under normal closed-loop operation with limited input excitation, the identified model captures only a restricted range of system dynamics. This directly affects SPC's prediction mechanism, where strict equality constraints require the Hankel matrix to be well-conditioned—a condition that is difficult to satisfy with poorly excited data. Additionally, subspace identification requires an observer design step to obtain reliable initial condition estimates, which adds to the overall tuning effort. Despite these modeling challenges, the identified state-space model remains linear and of small matrix size, so the online MPC optimization stays computationally light. As observed in the Pareto front comparisons, the modeling limitations cause degraded closed-loop performance in the comfort-consumption trade-off. SPC represents an interesting balance between data requirements and algorithmic complexity, making it suitable when moderate data availability and computational resources are available.

DeePC achieves the best control optimality among the three approaches. Under identical energy consumption levels, DeePC reduces thermal discomfort by approximately 1.2 $^{\circ}\mathrm{C} \cdot \mathrm{hour}$ compared to RC-MPC, and reaches comfort levels that are beyond the reach of both RC-MPC and SPC. This performance stems from DeePC's ability to optimize the underlying data-driven model and initial condition estimation online at each time step, rather than relying on a fixed offline-identified model. However, this comes at a cost: the Hankel matrix used for trajectory prediction grows with dataset length and estimation order, which significantly increases memory usage and places a heavy burden on the solver even though the optimization problem remains convex. This memory constraint in practice limits the admissible choices of initial condition estimation order T_{ini} and dataset length, restricting the richness of the dynamic representation. As a result, DeePC is best suited for systems with naturally low or reducible dynamic order. Within this scope, DeePC offers a very convenient

path to fast predictive control development, as no explicit modeling or system identification step is required. The main tuning effort lies in selecting the regularization weights λ_y and λ_g and the order T_{ini}, which require careful sensitivity analysis as demonstrated in Chapter 3. DeePC represents a theoretically elegant approach that is sensitive to data quality and computationally expensive, best deployed when abundant high-quality data and high computational capacity are available.

1.3 Adaptive Mechanisms

The thesis further extends both parametric and model-free predictive control to adaptive implementations, addressing time-varying dynamics induced by seasonal weather variations and equipment aging. Static approaches rely on a fixed dataset collected before deployment, which becomes progressively less representative as operating conditions shift. Chapter 4 demonstrates that this mismatch leads to measurable prediction accuracy degradation, motivating online update mechanisms.

Two update paradigms are investigated for the Hankel matrix used in DeePC. The **sliding window** approach maintains a fixed matrix dimension by replacing the oldest data column with each new measurement. This strategy achieves asymptotic adaptation after a complete window renewal, but produces sharp transient prediction error peaks immediately following abrupt parameter changes, as incoming columns temporarily mix pre- and post-shift dynamics. The **recursive reduction** approach compresses the growing dataset through fast SVD updates, preserving essential information in a low-dimensional subspace. This method avoids abrupt error peaks but retains a persistent residual error without explicit forgetting, since historical data from the previous operating regime remains present in the compressed basis. A forgetting factor $\gamma \in (0, 1]$ is introduced to bridge these two paradigms: exponential temporal weighting decays the influence of older measurements smoothly, enabling shorter effective memory lengths without violating the persistent excitation condition. For $\gamma = 0.99$, the effective memory spans approximately 300 samples, corresponding to about 12.5 days under the 15-minute sampling rate used in this study.

Three complementary indicators are combined in proposed strategy: the prediction error indicator I_{pred} monitors accuracy degradation in real time; the informativity indicator I_{info} evaluates persistent excitation through singular value rank analysis; and the operational relevance indicator I_{oper} prioritizes data collected during occupied service hours. This multi-criteria approach prevents a critical failure mode observed in baseline methods: when the system operates under quasi-constant inputs—common during comfort maintenance periods—new measurements are nearly collinear with existing columns, causing rank deficiency in the Hankel matrix and subsequent control degradation. The quality indicators detect and reject such measurements before they corrupt the data matrix.

Validation on the MOSA benchmark confirms that both adaptive DeePC variants achieve substantially superior Pareto fronts compared to static DeePC and SPC strategies, simultaneously reducing energy consumption and thermal comfort violations. An interesting structural result also emerges: the recursive reduction variant produces Pareto characteristics closely resembling those of adaptive SPC. Both methods construct low-dimensional projections of system dynamics—recursive reduction through SVD-based compression, adaptive SPC through recursive parametric state-space identification—leading to nearly identical energy-comfort trade-off curves despite their different mathematical formulations. The sliding window variant, by contrast, preserves raw measured trajectory columns without intermediate compression, producing a Pareto front more closely aligned with static DeePC behavior. These results highlight that the choice between update strategies has implications

beyond computational cost, influencing the nature of the dynamic representation and ultimately the control performance achievable in a given operating context.

1.4 Thesis Conclusion

This manuscript has addressed the modeling challenge in predictive control for building thermal management through data-driven methodologies. Three principal approaches have been examined: RC-MPC, SPC, and DeePC. These frameworks reduce dependency on first-principles modeling while preserving control performance.

Several aspects have been investigated to construct a complete control design solution. Hyperparameter analysis highlights their influence on closed-loop performance and computational tractability. Model extraction methodologies identify system dynamics from operational data. Adaptive strategies accommodate time-varying system characteristics. Integration of predictive control within cascade architectures through VRFT maintains the model-free methodology at the tracking level.

The synthesis demonstrates that successful deployment requires balancing multiple considerations. When physical insight exists and data are scarce, physics-informed models provide robust performance. When abundant operational data exist without modeling expertise, model-free implementations eliminate identification procedures. When system characteristics evolve, adaptive mechanisms maintain long-term performance. The hierarchical framework combining model-free tracking with various supervisory approaches provides solutions for heterogeneous building applications.

1.5 Valorisation and Scientific Contributions

The methods and frameworks developed throughout this thesis have been implemented in the Python toolkit `hiercon`. This toolkit provides a comprehensive platform for data-driven predictive control in building thermal management, including model identification methods (RC-based and subspace approaches), predictive control algorithms (RC-MPC, SPC, and DeePC), hierarchical compilation tools, and co-simulation benchmarking capabilities. The toolkit enables practitioners to apply the proposed control methods to real building systems without extensive custom development, facilitating broader adoption of data-driven predictive control in the HVAC industry.

The research results have been disseminated through peer-reviewed publications in leading journals and conference proceedings (see list in chapter I section 1.3). These publications cover the hierarchical control framework, comparative analysis of predictive control approaches, adaptive mechanisms for time-varying dynamics, and theoretical foundations for data-driven control in building thermal management. The publications contribute to the scientific community's understanding of practical data-driven control deployment in buildings.

Journal

- Building thermal control: Hierarchical design from limited data using gray-box or black-box internal models for model predictive control. *Journal of Building Engineering*, 2025.

Conference

- Hierarchical Model Predictive Control for Building Thermal Management: A mixed model-based and data-based approach. In *Proc. 11th International Conference on Control, Decision and Information Technologies (CoDIT)*, Split, Croatia, 2025. IEEE.
- Hierarchical Data-enabled Predictive Control for Building Thermal Management. *IFAC Control Application of Optimization Workshop*, 2025.

Paper under preparation

- Yuqi Liu, Pauline Kergus, Fabien Claveau, and Philippe Chevrel. Online Adaptation of Data-Enabled Predictive Control: Pro and Cons for Optimized Building HVAC Management. En préparation.

Software implementation

- HierCon: A Python Toolkit for Hierarchical and Optimized Thermal Management in Buildings. Internal Report, IMT Atlantique, 2025.

2 Perspectives

2.1 Perspectives on Data-Driven Predictive Control

Data-enabled predictive control presents two characteristics. The first is the use of column space from behavioral data to construct prediction trajectories. This concept has a link in depth with SPC, which has employed this principle and demonstrated its validity through extensive applications. The underlying mechanism relies on subspace projection to extract system dynamics from input-output data, establishing predictions through linear combinations of past trajectories.

The second characteristic distinguishes DeePC from SPC: dynamic model composition during each control step. Rather than extracting a fixed model offline, DeePC performs online linear combination among column vectors to construct the underlying model at each control iteration. This dynamic composition enables superior handling of nonlinearity by selecting preferred columns through regularization. The regularization mechanism effectively weights the composition of different LTI system snapshots according to their relevance to current operating conditions. Chapter III demonstrates this advantage through improved prediction accuracy compared to SPC.

However, this flexibility introduces tuning complexity. The regularization terms augment the cost function with additional weighting parameters that balance two competing objectives: fidelity to historical behavioral snapshots versus flexibility in linear composition. These regularization weights quantify the permitted deviation between the Hankel matrix representation and online measurements. As established in chapter III, these weights regularize the distance from system behavior embedded in the Hankel matrix to current operational data. Consequently, when economic objective weights vary to reflect different operational priorities, prediction performance fluctuates accordingly. This coupling between economic objective weights and regularization parameters creates additional tuning burden and demands elevated expertise in practical applications.

The practical application of DDPC approaches encounters computational challenges related to Hankel matrix dimensions. The row count depends on the initial state regression order and prediction horizon, while the column count represents the number of behavioral snapshots retained. Complex nonlinearities and extended prediction horizons substantially increase matrix dimensions, making computational burden a critical limitation of DeePC. Promising directions to address this include Koopman operator theory for nonlinear lifting or frequency-domain formulations for prediction and regression. In both cases, Willems' lemma based predictors still perform a linear composition, but no longer over raw system snapshots, rather over a lifted space [158] or an orthogonal basis. This reduces the matrix dimensions entering the optimization problem.

2.2 Perspectives on Building Control Implementation

The hierarchical data-driven predictive control framework proposed in this thesis has been developed and evaluated through numerical benchmarks. The BOPTEST framework provides standardized simulation environments enabling reproducible comparative studies. However, several aspects require further investigation for real-world deployment.

A theoretical consideration for building thermal systems is their linear parameter-varying (LPV) behavior. Building thermal dynamics are inherently nonlinear, yet they can be locally represented through LPV frameworks where scheduling variables capture time-varying operating conditions such as weather and equipment aging. The BOPTEST benchmark exhibits this LPV character through varying weather conditions and occupancy patterns. Extending data-driven predictive control approaches to explicitly account for LPV dynamics presents both theoretical and practical opportunities. Chapter IV develops adaptive mechanisms that implicitly handle parameter variations through online data updates. Appendix B provides detailed theoretical analysis comparing how SPC and DeePC perform under parameter-varying conditions. This analysis reveals specific challenges for Willems' lemma-based predictors when current operating parameters differ from those represented in historical data. Future work investigating explicit LPV extensions of DeePC, informed by the theoretical framework in appendix B, would strengthen the applicability of these methods to building systems experiencing significant seasonal or long-term performance changes.

The benchmark simulations operate under noise-free conditions, consequently neither robust formulations nor state filtering designs are incorporated in this work. Practical building systems exhibit measurement noise and unmeasured disturbances that degrade control performance. Recent developments in robust DeePC incorporate noise treatment directly into the Hankel formulation, addressing system input disturbances through structured regularization [112]. Evaluating robust DeePC performance within hierarchical control architectures presents particular interest. The tracking control error at the lower level propagates as input disturbance to the supervisory thermal power system. Characterizing how robust formulations mitigate this error propagation would establish practical feasibility.

Additional practical considerations include data quality management during continuous operation. The adaptive mechanisms investigated in chapter IV provide foundational strategies for maintaining prediction quality under time-varying dynamics. However, frameworks for detecting degraded data quality, triggering re-identification, and managing transient performance during adaptation require further development. The sliding data window and recursive reduction paradigms demonstrate

feasibility, yet comprehensive guidelines for threshold selection and adaptation rates remain application-specific.

Integration with existing building management systems constitutes another implementation challenge. The proposed framework requires access to operational data, computational resources for optimization, and communication infrastructure for hierarchical coordination. Deployment strategies that accommodate legacy control systems while progressively incorporating data-driven predictive capabilities warrant investigation. Hybrid approaches combining existing rule-based supervisory control with model-free tracking layers may provide practical transition pathways.

Beyond the challenges addressed in this thesis, several directions would further extend the practical value of the proposed framework. The Hankel matrix structure offers potential for fault detection in building systems. Deviations between predicted and measured trajectories that persist despite adaptive updates may indicate sensor failures or equipment faults rather than model drift. Residual-based fault isolation methods exploiting the subspace structure of the Hankel matrix have shown potential for this purpose [159, 160]. In the context of this work, the data quality indicators developed in Chapter 4 already monitor prediction residuals, providing a tool from which fault detection logic could be derived.

The current framework treats weather variables as measured disturbances provided by perfect forecasts. In practice, forecast uncertainty degrades control performance, particularly over long prediction horizons. Embedding probabilistic weather predictions into the supervisory optimization, as investigated in stochastic MPC formulations [8, 9], would improve robustness under realistic forecast errors. For data-driven approaches, uncertain disturbance trajectories could be incorporated through scenario-based extensions of the DeePC formulation [161].

Scaling the proposed hierarchical framework beyond individual buildings toward community-level or microgrid applications is also a interesting direction [162]. Coordinating multiple buildings sharing a common energy network introduces coupling constraints and grid-level objectives such as peak shaving or demand response. The data-driven nature of the proposed methods, which avoids building-specific physical models, facilitates replication across heterogeneous building stock, making community-scale deployment more tractable.

Appendix A

Extension to 2nd order RC model

The second-order RC model provides a more detailed representation of building thermal dynamics by separating the external wall thermal mass from the zone air thermal mass. This structure introduces additional states, capturing the thermal inertia of the building envelope. The model distinguishes between the zone air temperature $T_{z,i}$ and the external wall temperature $T_{wex,i}$, with corresponding thermal capacitances $C_{z,i}$ and $C_{wex,i}$. Three thermal resistances characterize the heat transfer paths: $R_{i,ex}$ represents the resistance between the zone air and the wall inner surface, $R_{ex,i}$ represents the resistance from the wall to the exterior environment, and $R_{tb,i}$ represents a thermal bridge connecting the zone air directly to the exterior, bypassing the wall mass.

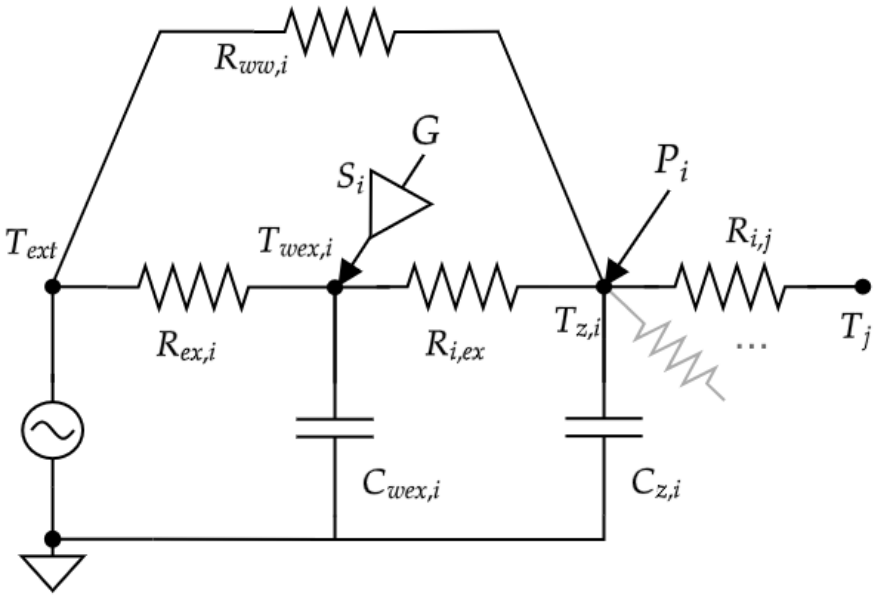

Figure A.1: Representation of a zone in 2nd order single zone RC model.

The equations for a single zone in the second-order model are:

$$C_{z,i}\dot{T}_{z,i} = \sum_{\forall j \in N_{ad,i}} \frac{1}{R_{i,j}}(T_{z,j} - T_{z,i}) + \frac{1}{R_{i,ex}}(T_{wex,i} - T_{z,i}) + \frac{1}{R_{tb,i}}(T_{ext} - T_{z,i}) + P_i \quad \text{(A.1a)}$$

$$C_{wex,i}\dot{T}_{wex,i} = \frac{1}{R_{i,ex}}(T_{z,i} - T_{wex,i}) + \frac{1}{R_{ex,i}}(T_{ext} - T_{wex,i}) + S_i G \quad \text{(A.1b)}$$

where $T_{z,i}$ is the air temperature in zone i and $T_{wex,i}$ is the wall temperature in zone i. The heating power P_i acts on the air node, while the solar radiation $S_i G$ is absorbed by the wall node. For a single zone i, the state-space model can be expressed as:

$$\begin{bmatrix} \dot{T}_{z,i} \\ \dot{T}_{wex,i} \end{bmatrix} = A_1(\theta_{phy,i}) \begin{bmatrix} T_{z,i} \\ T_{wex,i} \end{bmatrix} + A_2(\theta_{phy,i}) \begin{bmatrix} T_{z,j_1} \\ \vdots \\ T_{z,j_m} \end{bmatrix} + B_1(\theta_{phy,i}) P_i + B_2(\theta_{phy,i}) \begin{bmatrix} G \\ T_{ext} \end{bmatrix} \tag{A.2a}$$

$$A_1(\theta_{phy,i}) = \begin{bmatrix} -\left(\frac{1}{R_{i,ex}} + \frac{1}{R_{tb,i}} + \sum_{j \in N_{ad,i}} \frac{1}{R_{i,j}}\right) \frac{1}{C_{z,i}} & \frac{1}{R_{i,ex} C_{z,i}} \\ \frac{1}{R_{i,ex} C_{wex,i}} & -\left(\frac{1}{R_{i,ex}} + \frac{1}{R_{ex,i}}\right) \frac{1}{C_{wex,i}} \end{bmatrix} \tag{A.2b}$$

$$A_2(\theta_{phy,i}) = \begin{bmatrix} \frac{1}{R_{i,j_1} C_{z,i}} & \cdots & \frac{1}{R_{i,j_m} C_{z,i}} \\ 0 & \cdots & 0 \end{bmatrix} \tag{A.2c}$$

$$B_1(\theta_{phy,i}) = \begin{bmatrix} \frac{1}{C_{z,i}} \\ 0 \end{bmatrix} \tag{A.2d}$$

$$B_2(\theta_{phy,i}) = \begin{bmatrix} 0 & \frac{1}{R_{tb,i} C_{z,i}} \\ \frac{S_i}{C_{wex,i}} & \frac{1}{R_{ex,i} C_{wex,i}} \end{bmatrix} \tag{A.2e}$$

In the following, we denote elements of A and B matrices using superscript notation: $A_k^{(a,b)}(\theta_{phy,i})$ or $B_k^{(a,b)}(\theta_{phy,i})$, where subscript $k \in \{1, 2\}$ identifies which matrix (A_1 or A_2, B_1 or B_2), superscript (a, b) denotes the row a and column b position within that matrix, and i refers to the zone index. For example, $A_1^{(1,2)}(\theta_{phy,i}) = \frac{1}{R_{i,ex} C_{z,i}}$ is the element at row 1, column 2 of the matrix A_1 for zone i.

The physical parameter vector for zone i is:

$$\theta_{phy,i} = [R_{i,ex},\ R_{ex,i},\ R_{tb,i},\ C_{z,i},\ C_{wex,i},\ S_i,\ R_{i,j_1},\ R_{i,j_2},\ \ldots,\ R_{i,j_m}]^\top \tag{A.3}$$

where $j_1, j_2, \ldots, j_m \in N_{ad,i}$ are the indices of zones adjacent to zone i. The full multi-zone physical parameter vector is:

$$\boldsymbol{\theta}_{phy} = \left[\theta_{phy,1}^\top,\ \theta_{phy,2}^\top,\ \ldots,\ \theta_{phy,n}^\top\right]^\top \tag{A.4}$$

To extend this to a multi-zone model, we concatenate the state vectors and arrange the matrices accordingly:

$$\begin{bmatrix} \dot{T}_{z,1} \\ \vdots \\ \dot{T}_{z,n} \\ \dot{T}_{wex,1} \\ \vdots \\ \dot{T}_{wex,n} \end{bmatrix} = \mathbf{A}(\boldsymbol{\theta}_{phy}) \begin{bmatrix} T_{z,1} \\ \vdots \\ T_{z,n} \\ T_{wex,1} \\ \vdots \\ T_{wex,n} \end{bmatrix} + \mathbf{B}_1(\boldsymbol{\theta}_{phy}) \begin{bmatrix} P_1 \\ \vdots \\ P_n \end{bmatrix} + \mathbf{B}_2(\boldsymbol{\theta}_{phy}) \begin{bmatrix} G \\ T_{ext} \end{bmatrix} \tag{A.5}$$

The system matrix $\mathbf{A}(\boldsymbol{\theta}_{phy})$ of size $(2n \times 2n)$ has the following block structure:

$$\mathbf{A}(\boldsymbol{\theta}_{phy}) = \begin{bmatrix} \mathbf{A}^{zz} & \mathbf{A}^{zw} \\ \mathbf{A}^{wz} & \mathbf{A}^{ww} \end{bmatrix} \tag{A.6}$$

where each block is of size $(n \times n)$:

- $\mathbf{A}^{zz}$: air-to-air coupling. The diagonal entry (i, i) is the self-dynamics of the air node in zone i. The off-diagonal entry (i, j) is non-zero only if zone j is adjacent to zone i ($j \in N_{ad,i}$), representing inter-zone heat exchange through the shared resistance $R_{i,j}$.

$$\mathbf{A}^{zz} = \begin{bmatrix} A_1^{(1,1)}(\theta_{phy,1}) & A_2^{(1,1)}(\theta_{phy,1}) & \cdots & A_2^{(1,m_1)}(\theta_{phy,1}) \\ A_2^{(1,1)}(\theta_{phy,2}) & A_1^{(1,1)}(\theta_{phy,2}) & \cdots & A_2^{(1,m_2)}(\theta_{phy,2}) \\ \vdots & \vdots & \ddots & \vdots \\ A_2^{(1,1)}(\theta_{phy,n}) & A_2^{(1,2)}(\theta_{phy,n}) & \cdots & A_1^{(1,1)}(\theta_{phy,n}) \end{bmatrix} \tag{A.7}$$

- $\mathbf{A}^{zw}$: wall-to-air coupling. A diagonal matrix where entry (i, i) is $A_1^{(1,2)}(\theta_{phy,i})$, representing the heat flow from the wall to the air in zone i through the internal wall resistance $R_{i,ex}$.

$$\mathbf{A}^{zw} = \begin{bmatrix} A_1^{(1,2)}(\theta_{phy,1}) & 0 & \cdots & 0 \\ 0 & A_1^{(1,2)}(\theta_{phy,2}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_1^{(1,2)}(\theta_{phy,n}) \end{bmatrix} \tag{A.8}$$

- $\mathbf{A}^{wz}$: air-to-wall coupling. A diagonal matrix where entry (i, i) is $A_1^{(2,1)}(\theta_{phy,i})$, representing the heat flow from the air to the wall in zone i through the internal wall resistance $R_{i,ex}$.

$$\mathbf{A}^{wz} = \begin{bmatrix} A_1^{(2,1)}(\theta_{phy,1}) & 0 & \cdots & 0 \\ 0 & A_1^{(2,1)}(\theta_{phy,2}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_1^{(2,1)}(\theta_{phy,n}) \end{bmatrix} \tag{A.9}$$

- $\mathbf{A}^{ww}$: wall self-dynamics. A diagonal matrix where entry (i, i) is $A_1^{(2,2)}(\theta_{phy,i})$, representing the thermal losses of the wall node in zone i toward both the air and the exterior.

$$\mathbf{A}^{ww} = \begin{bmatrix} A_1^{(2,2)}(\theta_{phy,1}) & 0 & \cdots & 0 \\ 0 & A_1^{(2,2)}(\theta_{phy,2}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_1^{(2,2)}(\theta_{phy,n}) \end{bmatrix} \tag{A.10}$$

The input and disturbance matrices are:

$$\mathbf{B}_1(\boldsymbol{\theta}_{phy}) = \begin{bmatrix} B_1^{(1,1)}(\theta_{phy,1}) & 0 & \cdots & 0 \\ 0 & B_1^{(1,1)}(\theta_{phy,2}) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & B_1^{(1,1)}(\theta_{phy,n}) \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{bmatrix} \tag{A.11a}$$

$$\mathbf{B}_2(\boldsymbol{\theta}_{phy}) = \begin{bmatrix} B_2^{(1,1)}(\theta_{phy,1}) & B_2^{(1,2)}(\theta_{phy,1}) \\ \vdots & \vdots \\ B_2^{(1,1)}(\theta_{phy,n}) & B_2^{(1,2)}(\theta_{phy,n}) \\ B_2^{(2,1)}(\theta_{phy,1}) & B_2^{(2,2)}(\theta_{phy,1}) \\ \vdots & \vdots \\ B_2^{(2,1)}(\theta_{phy,n}) & B_2^{(2,2)}(\theta_{phy,n}) \end{bmatrix} \tag{A.11b}$$

As in the first-order case, the physical parameters $\boldsymbol{\theta}_{phy}$ enter the matrices through products of resistances and capacitances, making the model nonlinear in $\boldsymbol{\theta}_{phy}$. We introduce an intermediate parameter vector $\boldsymbol{\theta}_{opt}$ that is linear in the data columns. For the air node of zone i:

$$\theta_{opt,z,i} = \Bigg[-\left(\frac{1}{R_{i,ex}} + \frac{1}{R_{tb,i}} + \sum_{j \in N_{ad,i}} \frac{1}{R_{i,j}} \right) \frac{1}{C_{z,i}},\ \frac{1}{R_{i,j_1} C_{z,i}},\ \ldots,\ \frac{1}{R_{i,j_m} C_{z,i}}, \\ \frac{1}{R_{i,ex} C_{z,i}},\ \frac{1}{C_{z,i}},\ \frac{1}{R_{tb,i} C_{z,i}} \Bigg]^\top \tag{A.12}$$

For the wall node of zone i:

$$\theta_{opt,w,i} = \left[-\left(\frac{1}{R_{i,ex}} + \frac{1}{R_{ex,i}} \right) \frac{1}{C_{wex,i}},\ \frac{1}{R_{i,ex} C_{wex,i}},\ \frac{S_i}{C_{wex,i}},\ \frac{1}{R_{ex,i} C_{wex,i}} \right]^\top \tag{A.13}$$

The full intermediate parameter vector is:

$$\boldsymbol{\theta}_{opt} = \left[\theta_{opt,z,1}^\top,\ \theta_{opt,w,1}^\top,\ \ldots,\ \theta_{opt,z,n}^\top,\ \theta_{opt,w,n}^\top \right]^\top \tag{A.14}$$

Observability analysis

For the second-order single-zone model with state vector $[T_{z,i},\ T_{wex,i}]^\top$ and output matrix $C = [1 \quad 0]$ (only air temperature is measured), the observability matrix is:

$$\mathcal{O}_i = \begin{bmatrix} C \\ CA_1(\theta_{phy,i}) \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ A_1^{(1,1)}(\theta_{phy,i}) & A_1^{(1,2)}(\theta_{phy,i}) \end{bmatrix} \tag{A.15}$$

The system is observable if $\mathcal{O}_i$ has full rank 2. Since $A_1^{(1,2)}(\theta_{phy,i}) = \frac{1}{R_{i,ex} C_{z,i}} \neq 0$ for non-zero resistance and capacitance, the matrix has full rank. This means that the wall temperature $T_{wex,i}$ can be reconstructed from measurements of the air temperature $T_{z,i}$ alone, given the system dynamics.

For the multi-zone system with output matrix $\mathbf{C} = [I_n \quad 0_n]$ (measuring all n zone air temperatures), the observability matrix is:

$$\mathcal{O} = \begin{bmatrix} \mathbf{C} \\ \mathbf{CA}(\boldsymbol{\theta}_{phy}) \\ \mathbf{CA}^2(\boldsymbol{\theta}_{phy}) \\ \vdots \\ \mathbf{CA}^{2n-1}(\boldsymbol{\theta}_{phy}) \end{bmatrix} \tag{A.16}$$

The system is observable if $\mathcal{O}$ has full rank $2n$. This condition is satisfied when all thermal resistances and capacitances are non-zero and finite, inter-zone couplings are distinct, and the thermal network contains no perfect symmetries. The first n rows directly observe the zone air temperatures $T_{z,i}$. The next n rows from $\mathbf{CA}(\boldsymbol{\theta}_{phy})$ couple air and wall temperatures through the non-zero thermal resistances, providing information about $T_{wex,i}$. All internal states, including wall temperatures, can therefore be reconstructed from zone air temperature measurements alone.

Challenges in Modeling

Higher-order models require dual estimation. One of the primary challenges in applying least square estimation to RC models is the unavailability of certain state measurements, particularly wall temperatures. Higher-order RC models require temperature measurements on additional thermal capacitances, such as walls. This leads to a dual estimation problem where both states and parameters need to be estimated simultaneously. The dual estimation problem is inherently nonlinear, as it involves products of unknown states and parameters. For example, the 2nd order model in equation (A.1). The term $\frac{1}{R_{ex}} T_{wex}$ involves a product of an unknown parameter ($\frac{1}{R_{ex}}$) and an unknown state (T_{wex}).

The alternative to dual estimation is to estimate unmeasured temperatures as system states using methods like extended Kalman filtering. However, the core difficulty is that these estimated states are not physical states, but their projections onto the reduced-order data-driven model. Since the model order is significantly reduced and data conditions are not ideal, these states can deviate substantially from physical reality. This deviation leads to difficulties in observer design and tuning, as the estimated states lose their physical interpretability and the gray-box model effectively becomes partially black-boxed.

Appendix B

SPC and DDPC comparison for LPV systems

The static versus dynamic modeling paradigms lead to different behaviors under parameter variations. SPC optimizes the input trajectory through a reduced-order system projection. The data is summarized as a parametric model in advance through (80), which represents an averaged system behavior across the dataset. When parameters vary, this averaged model becomes increasingly inaccurate as it cannot capture the full complexity of parameter-dependent dynamics.

DDPC, in contrast, performs dynamic modeling with input composition. The vector g is recomputed at each control step, allowing different combinations of historical snapshots depending on current conditions and control objectives. This difference emerges distinctly in two cases:

Training Set Validation

Consider validating the identified models on the training dataset itself. For SPC, we predict outputs using:

$$\hat{\mathbf{y}}_i = \bar{A}^*_{ini}\mathbf{y}_{p,i} + \bar{B}^*_{ini}\mathbf{u}_{p,i} + \bar{B}^*_f\mathbf{u}_{f,i}, \tag{B.1}$$

where subscript i indexes the i-th trajectory in the dataset. The prediction error is:

$$e_{SPC} = \sum_{i=1}^{T_d - T_{ini} - N + 1} ||\mathbf{y}_{f,i} - \hat{\mathbf{y}}_i||_2^2. \tag{B.2}$$

If the system is not truly LTI, subspace identification will always yield a residual error $e_{SPC} > 0$. The reduced-order projection cannot perfectly represent all trajectories simultaneously when system parameters vary across the dataset. This residual reflects the approximation error inherent in summarizing diverse behaviors into a single parametric model.

For DDPC validation on the same dataset, we solve:

$$g_i^* = \arg\min_g \left\| \begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g - \begin{bmatrix} \mathbf{u}_{p,i} \\ \mathbf{y}_{p,i} \\ \mathbf{u}_{f,i} \\ \mathbf{y}_{f,i} \end{bmatrix} \right\|_2^2. \tag{B.3}$$

The optimal solution is a canonical vector:

$$g_i^* = e_i = [0, \ldots, 0, \underbrace{1}_{i\text{-th position}}, 0, \ldots, 0]^\top, \tag{B.4}$$

which selects exactly the i-th column of the Hankel matrices. This gives zero prediction error:

$$e_{DDPC} = \sum_{i=1}^{T_d - T_{ini} - N + 1} ||\mathbf{y}_{f,i} - Y_f g_i^*||_2^2 = 0. \tag{B.5}$$

This perfect reconstruction capability demonstrates that DDPC preserves all system behavior snapshots without approximation. Each trajectory remains directly accessible through appropriate selection of g. In contrast, SPC's reduced-order projection inherently discards information, leading to unavoidable approximation errors when system behavior varies.

Control objective as regularization

The control objective $J(\mathbf{y}, \mathbf{u})$ in DDPC formulations acts as task-specific regularization that guides which historical data receives emphasis during prediction.

Consider the standard DeePC optimization problem:

$$\min_{g,\sigma_y} \quad J(\mathbf{y}, \mathbf{u}) + \lambda_y ||\sigma_y||_2^2 + \lambda_g ||g||_p \tag{B.6a}$$

$$\text{s.t.} \quad \begin{bmatrix} U_p \\ Y_p \\ U_f \\ Y_f \end{bmatrix} g = \begin{bmatrix} \mathbf{u}_{ini} \\ \mathbf{y}_{ini} \\ \mathbf{u} \\ \mathbf{y} \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma_y \\ 0 \\ 0 \end{bmatrix}, \tag{B.6b}$$

$$\mathbf{u} \in \mathcal{U}^N, \quad \mathbf{y} \in \mathcal{Y}^N, \tag{B.6c}$$

where $J(\mathbf{y}, \mathbf{u})$ typically penalizes tracking error and control effort, λ_y and λ_g are regularization weights, and $\mathcal{U}^N$ and $\mathcal{Y}^N$ denote constraint sets for inputs and outputs over the prediction horizon.

The regularization term $\lambda_g ||g||_p$ (commonly with $p \in \{1, 2\}$) encourages sparsity or smoothness in the trajectory combination vector g. This regularization implicitly selects historical data that aligns with current operating conditions. When current parameter values $\theta[t_k]$ resemble those in certain dataset segments, the corresponding columns in the Hankel matrices yield better tracking performance in $J(\mathbf{y}, \mathbf{u})$ with smaller σ_y, leading to larger weights in g for those columns.

Conversely, historical data collected under dissimilar parameter conditions contributes less to the optimal g. If forced to rely heavily on such data, the constraint (B.6b) cannot be satisfied tightly, resulting in large σ_y that incurs penalty through $\lambda_y ||\sigma_y||_2^2$. The optimization naturally reduces contributions from irrelevant historical snapshots, effectively performing online model selection based on local relevance. The interplay between control objectives and dataset selection will be further developed in subsequent sections on adaptive DDPC.

Data quality

The persistency of excitation (PE) condition in Definition 1 ensures that the input signal explores all directions in the input space of an LTI system. For LTI systems, this condition guarantees that the Hankel matrix has full row rank, enabling complete characterization of system dynamics. The PE condition verifies that inputs excite all system modes up to order L, which suffices for LTI subspace identification and data-driven control.

For LPV systems described by (86), the PE condition alone becomes insufficient. While a persistently exciting input explores the input space $\mathbb{R}^m$, it does not guarantee exploration of the

parameter space. The scheduling parameter trajectory $\{\theta[k]\}_{k=0}^{T_d-1}$ often depends on uncontrollable external conditions, making it infeasible to satisfy PE requirements through deliberate excitation. Consider a building thermal system where scheduling parameters depend on weather conditions—such as heat transfer coefficients varying with solar radiation and wind speed. Even with highly exciting control inputs for heating and cooling, the parameter trajectory remains constrained by naturally occurring weather patterns. Data collected during predominantly cloudy winter months captures limited variation in solar-dependent parameters. The resulting dataset contains rich controllable input variations but sparse parameter coverage. When sunny summer conditions arrive, the pseudo-matrices $\bar{\mathcal{T}}(t_k)$ and $\bar{\mathcal{O}}(t_k)$ differ substantially from their cloudy-weather counterparts embedded in the dataset, leading to prediction errors.

Dataset richness for LPV systems thus requires dual conditions: (1) input space coverage through persistently exciting controllable signals satisfying Definition 1, and (2) parameter space coverage ensuring the scheduling trajectory spans representative operating conditions. The second condition poses challenges since scheduling parameters often depend on uncontrollable environmental factors that cannot be artificially excited. Closed-loop operational data compounds these difficulties—regulatory control action reduces controllable input excitation, while stable environmental periods limit natural parameter variation. These challenges motivate adaptive mechanisms that continuously update datasets while maintaining both input and parameter richness.

The DeePC framework developed in Chapter III relies on a static data matrix constructed from offline identification experiments. This static approach assumes that system dynamics remain constant throughout operation, which rarely holds for building thermal systems. Buildings experience continuous changes in thermal behavior due to seasonal weather variations, occupancy pattern shifts, equipment aging, and evolving operational strategies. These time-varying characteristics require control methods that can track dynamic evolution without periodic reidentification experiments.

This section develops adaptive mechanisms that update the data matrix online according to real-time measurements. The adaptive framework addresses two challenges: maintaining prediction accuracy as system dynamics evolve, and managing computational complexity as new data accumulates. The development follows a progression. First, the updating mechanism subsection establishes how to incorporate new measurements into the data matrix through sliding window and recursive reduction paradigms, then introduces forgetting mechanisms that assign temporal weights to prioritize recent information. Second, the data quality assignment subsection develops criteria for evaluating whether new measurements should be incorporated, preventing degradation from informationally deficient data.

References

[1] D. Blum, J. Arroyo, S. Huang, J. Drgoňa, F. Jorissen, H. T. Walnum, Y. Chen, K. Benne, D. Vrabie, M. Wetter et al., "Building optimization testing framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings," Journal of Building Performance Simulation, vol. 14, no. 5, pp. 586–610, 2021.

[2] Z. Wang and Y. Chen, "Data-driven modeling of building thermal dynamics: Methodology and state of the art," Energy and Buildings, vol. 203, p. 109405, 2019.

[3] " World Energy Outlook 2024," IEA (2024). [Online]. Available: https://www.iea.org/reports/world-energy-outlook-2024

[4] U. S. S. O. Engineering, "Phase Change Composite Materials for Energy Efficient Building Envelopes," UCLA Samueli School Of Engineering, Tech. Rep., 2018. [Online]. Available: https://www.seas.ucla.edu/~pilon/PCMIntro.html

[5] E. P. of Buildings Directive, "DIRECTIVE 2024/1275 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 24 April 2024 on the energy performance of buildings," 2024.

[6] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, "All you need to know about model predictive control for buildings," Annual Reviews in Control, vol. 50, pp. 190–232, 2020.

[7] K. Bamdad, N. Mohammadzadeh, M. Cholette, and S. Perera, "Model Predictive Control for Energy Optimization of HVAC Systems Using EnergyPlus and ACO Algorithm," Buildings, vol. 13, no. 12, p. 3084, 2023.

[8] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," Energy and Buildings, vol. 45, pp. 15–27, 2012.

[9] F. Oldewurtel, C. N. Jones, A. Parisio, and M. Morari, "Stochastic Model Predictive Control for Building Climate Control," IEEE Transactions on Control Systems Technology, vol. 22, no. 3, pp. 1198–1205, 2014.

[10] D. Sturzenegger, D. Gyalistras, M. Gwerder, C. Sagerschnig, M. Morari, and R. Smith, "Model Predictive Control of a Swiss office building," Clima 2013, pp. 3227–3236, 2013.

[11] Y. Ma, F. Borrelli, B. Hencey, B. Coffey, S. Bengea, and P. Haves, "Model predictive control for the operation of building cooling systems," IEEE Transactions on Control Systems Technology, vol. 20, no. 3, pp. 796–803, 2011.

[12] J. Arroyo, F. Spiessens, and L. Helsen, "Comparison of model complexities in optimal control tested in a real thermally activated building system," Buildings, vol. 12, no. 5, p. 539, 2022.

[13] F. Jorissen, D. Picard, K. Six, and L. Helsen, "Detailed white-box non-linear model predictive control for scalable building HVAC control," in Modelica Conferences, 2021, pp. 315–323.

[14] J. Drgoňa, D. Picard, and L. Helsen, "Cloud-based implementation of white-box model predictive control for a GEOTABS office building: A field test demonstration," Journal of Process Control, vol. 88, pp. 63–77, 2020.

[15] F. Gauthier-Clerc, H. Le Capitaine, F. Claveau, and P. Chevrel, "Comparing Neural Network and Linear Models in Economic MPC: Insights from BOPTEST for Building Temperature Control," in ECC 2024: European Control Conference, 2024.

[16] P. Van Overschee and B. De Moor, Subspace identification for linear systems: Theory - Implementation - Applications. Springer Science & Business Media, 2012.

[17] S. Prívara, Z. Váňa, D. Gyalistras, J. Cigler, C. Sagerschnig, M. Morari, and L. Ferkl, "Modeling and identification of a large multi-zone office building," in 2011 IEEE International Conference on Control Applications (CCA). IEEE, 2011, pp. 55–60.

[18] H. Huang, L. Chen, and E. Hu, "A neural network-based multi-zone modelling approach for predictive control system design in commercial buildings," Energy and Buildings, vol. 97, pp. 86–97, 2015.

[19] H. Y. Noh and R. Rajagopal, "Data-driven forecasting algorithms for building energy consumption," in Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2013, vol. 8692. SPIE, 2013, pp. 221–228.

[20] G. P. Henze, K. J. Kircher, and J. E. Braun, "Why has advanced commercial HVAC control not yet achieved its promise?" Journal of Building Performance Simulation, vol. 18, no. 2, pp. 217–228, 2025.

[21] X. Liu, M. Ren, Z. Yang, G. Yan, Y. Guo, L. Cheng, and C. Wu, "A multi-step predictive deep reinforcement learning algorithm for HVAC control systems in smart buildings," Energy, vol. 259, p. 124857, 2022.

[22] S. Zhou, A. Shah, P. Leung, X. Zhu, and Q. Liao, "A comprehensive review of the applications of machine learning for HVAC," DeCarbon, p. 100023, 2023.

[23] R. Yang and L. Wang, "Multi-zone building energy management using intelligent control and optimization," Sustainable Cities and Society, vol. 6, pp. 16–21, 2013.

[24] D. Zhao, D. Watari, Y. Ozawa, I. Taniguchi, T. Suzuki, Y. Shimoda, and T. Onoye, "Data-driven online energy management framework for HVAC systems: An experimental study," Applied Energy, vol. 352, p. 121921, 2023.

[25] S. Zhan and A. Chong, "Data requirements and performance evaluation of model predictive control in buildings: A modeling perspective," Renewable and Sustainable Energy Reviews, vol. 142, p. 110835, 2021.

[26] S. R. West, J. K. Ward, and J. Wall, "Trial results from a model predictive control and optimisation system for commercial building HVAC," Energy and Buildings, vol. 72, pp. 271–279, 2014.

[27] Y. Zong, G. M. Böning, R. M. Santos, S. You, J. Hu, and X. Han, "Challenges of implementing economic model predictive control strategy for buildings interacting with smart energy systems," Applied Thermal Engineering, vol. 114, pp. 1476–1486, 2017.

[28] A. Mirakhorli and B. Dong, "Market and behavior driven predictive energy management for residential buildings," Sustainable Cities and Society, vol. 38, pp. 723–735, 2018.

[29] J. Joe and P. Karava, "A model predictive control strategy to optimize the performance of radiant floor heating and cooling systems in office buildings," Applied Energy, vol. 245, pp. 65–77, 2019.

[30] L. Ljung, "System identification," in Signal Analysis and Prediction. Springer, 1998, pp. 163–173.

[31] P. Stoffel, L. Maier, A. Kümpel, T. Schreiber, and D. Müller, "Evaluation of advanced control strategies for building energy systems," Energy and Buildings, vol. 280, p. 112709, 2023.

[32] "Grey-box modeling and application for building energy simulations — A critical review," Renewable and Sustainable Energy Reviews, vol. 146, p. 111174, 2021.

[33] A. Lefort, R. Bourdais, G. Ansanay-Alex, and H. Guéguen, "Hierarchical control method applied to energy management of a residential house," Energy and Buildings, vol. 64, pp. 53–61, 2013.

[34] C. Vallianos, A. Athienitis, and B. Delcroix, "Automatic generation of multi-zone RC models using smart thermostat data from homes," Energy and Buildings, vol. 277, 2022.

[35] J. Arroyo, F. Spiessens, and L. Helsen, "Identification of multi-zone grey-box building models for use in model predictive control," Journal of Building Performance Simulation, vol. 13, no. 4, pp. 472–486, 2020.

[36] E. Biyik and A. Kahraman, "A predictive control strategy for optimal management of peak load, thermal comfort, energy storage and renewables in multi-zone buildings," Journal of Building Engineering, vol. 25, p. 100826, 2019.

[37] F. Langner, M. Frahm, W. Wang, J. Matthes, and V. Hagenmeyer, "Hierarchical-stochastic model predictive control for a grid-interactive multi-zone residential building with distributed energy resources," Journal of Building Engineering, vol. 89, p. 109401, 2024.

[38] Y. Zhang, Z. O'Neill, B. Dong, and G. Augenbroe, "Comparisons of inverse modeling approaches for predicting building energy performance," Building and Environment, vol. 86, pp. 177–190, 2015.

[39] Z. Afroz, G. Shafiullah, T. Urmee, and G. Higgins, "Modeling techniques used in building HVAC control systems: A review," Renewable and Sustainable Energy Reviews, vol. 83, pp. 64–84, 2018.

[40] M. C. Campi, A. Lecchini, and S. M. Savaresi, "Virtual reference feedback tuning: a direct method for the design of feedback controllers," Automatica, vol. 38, no. 8, pp. 1337–1346, 2002.

[41] J. G. Araque, L. Angel, J. Viola, and Y. Chen, "Design and implementation of a recursive feedforward-based virtual reference feedback tuning (VRFT) controller for temperature uniformity control applications," Machines, vol. 11, no. 10, p. 975, 2023.

[42] R. Scattolini, "Architectures for distributed and hierarchical model predictive control — a review," Journal of Process Control, vol. 19, no. 5, pp. 723–731, 2009.

[43] C. Wang, B. Wang, and F. You, "Demand response for residential buildings using hierarchical nonlinear model predictive control for plug-and-play," Applied Energy, vol. 369, p. 123581, 2024.

[44] J. C. Willems, P. Rapisarda, I. Markovsky, and B. L. De Moor, "A note on persistency of excitation," Systems & Control Letters, vol. 54, no. 4, pp. 325–329, 2005.

[45] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," in 2019 18th European Control Conference (ECC). IEEE, 2019, pp. 307–312.

[46] R. Baetens, R. De Coninck, F. Jorissen, D. Picard, L. Helsen, and D. Saelens, "Openideas — an open framework for integrated district energy simulations," in Building Simulation, 2015, pp. 345–354.

[47] M. Wetter, W. Zuo, T. S. Nouidui, and X. Pang, "Modelica buildings library," Journal of Building Performance Simulation, vol. 7, no. 4, pp. 253–270, 2014.

[48] U.S. Department of Energy, "ANSI/ASHRAE/IES Standard 90.1 Prototype Building Models," https://www.energycodes.gov/prototype-building-models.

[49] A. Abreu, R. Bourdais, and H. Guéguen, "Hierarchical Model Predictive Control for Building Energy Management of Hybrid Systems," IFAC-PapersOnLine, vol. 51, no. 16, pp. 235–240, 2018.

[50] E. Van Henten and J. Bontsema, "Time-scale decomposition of an optimal control problem in greenhouse climate management," Control Engineering Practice, vol. 17, no. 1, pp. 88–96, 2009.

[51] J. L. Proctor, S. L. Brunton, and J. N. Kutz, "Dynamic mode decomposition with control," SIAM Journal on Applied Dynamical Systems, vol. 15, no. 1, pp. 142–161, 2016.

[52] J. B. Rawlings, D. Q. Mayne, and M. Diehl, Model predictive control: theory, computation, and design. Nob Hill Publishing, 2020, vol. 2.

[53] J. Köhler, M. A. Müller, and F. Allgöwer, "Nonlinear reference tracking: An economic model predictive control perspective," IEEE Transactions on Automatic Control, vol. 64, no. 1, pp. 254–269, 2018.

[54] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," Journal of Process Control, vol. 24, no. 8, pp. 1156–1178, 2014.

[55] D. Angeli, R. Amrit, and J. B. Rawlings, "On average performance and stability of economic model predictive control," IEEE Transactions on Automatic Control, vol. 57, no. 7, pp. 1615–1626, 2011.

[56] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in 2012 IEEE 51st Conference on Decision and Control (CDC). IEEE, 2012, pp. 3851–3861.

[57] J. Drgoňa, M. Klaučo, and M. Kvasnica, "MPC-based reference governors for thermostatically controlled residential buildings," in 2015 54th IEEE Conference on Decision and Control (CDC). IEEE, 2015, pp. 1334–1339.

[58] V. Chinde, Y. Lin, and M. J. Ellis, "Data-enabled predictive control for building HVAC systems," Journal of Dynamic Systems, Measurement, and Control, vol. 144, no. 8, 2022.

[59] C. L. Remes, R. B. Gomes, J. V. Flores, F. B. Líbano, and L. Campestrini, "Virtual reference feedback tuning applied to DC–DC converters," IEEE Transactions on Industrial Electronics, vol. 68, no. 1, pp. 544–552, 2020.

[60] T. Sato, Y. Sakai, N. Kawaguchi, and O. Arrieta, "Dual-rate data-driven virtual reference feedback tuning: Improvement in fast-tracking performance and ripple-free design," IEEE Access, vol. 9, pp. 144 426–144 437, 2021.

[61] M. Benosman, "Model-based vs data-driven adaptive control: an overview," International Journal of Adaptive Control and Signal Processing, vol. 32, no. 5, pp. 753–776, 2018.

[62] P. Kergus, S. Formentin, M. Giuliani, and A. Castelletti, "Learning-based hierarchical control of water reservoir systems," IFAC Journal of Systems and Control, vol. 19, p. 100185, 2022.

[63] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," Advances in Neural Information Processing Systems, vol. 31, 2018.

[64] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," Journal of Computational Physics, vol. 378, pp. 686–707, 2019.

[65] L. Von Rueden, S. Mayer, K. Beckh, B. Georgiev, S. Giesselbach, R. Heese, B. Kirsch, J. Pfrommer, A. Pick, R. Ramamurthy et al., "Informed machine learning — a taxonomy and survey of integrating prior knowledge into learning systems," IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 1, pp. 614–633, 2021.

[66] T. Bohlin, Practical grey-box process identification: theory and applications. Springer, 2006.

[67] N. R. Kristensen, H. Madsen, and S. B. Jørgensen, "Parameter estimation in stochastic grey-box models," Automatica, vol. 40, no. 2, pp. 225–237, 2004.

[68] B. Ninness, "Estimation of 1/f noise," IEEE Transactions on Information Theory, vol. 44, no. 1, pp. 32–46, 2002.

[69] D. C. Psichogios and L. H. Ungar, "A hybrid neural network-first principles approach to process modeling," AIChE Journal, vol. 38, no. 10, pp. 1499–1511, 1992.

[70] A. Karpatne, W. Watkins, J. Read, and V. Kumar, "How can physics inform deep learning methods in scientific problems?: Recent progress and future prospects," in 31st Conference on Neural Information Processing Systems (NeurIPS), 2017, pp. 1–5.

[71] M. Hu and F. Xiao, "Price-responsive model-based optimal demand response control of inverter air conditioners using genetic algorithm," Applied Energy, vol. 219, pp. 151–164, 2018.

[72] S. F. Fux, A. Ashouri, M. J. Benz, and L. Guzzella, "EKF based self-adaptive thermal model for a passive house," Energy and Buildings, vol. 68, pp. 811–817, 2014.

[73] O. T. Ogunsola, L. Song, and G. Wang, "Development and validation of a time-series model for real-time thermal load estimation," Energy and Buildings, vol. 76, pp. 440–449, 2014.

[74] P. Van Overschee and B. De Moor, "Continuous-time frequency domain subspace system identification," Signal Processing, vol. 52, no. 2, pp. 179–194, 1996.

[75] C. M. Bishop and N. M. Nasrabadi, Pattern recognition and machine learning. Springer, 2006.

[76] C. K. Williams and C. E. Rasmussen, Gaussian processes for machine learning. MIT Press, 2006.

[77] Z. Wang, T. Hong, and M. A. Piette, "Building thermal load prediction through shallow machine learning and deep learning," Applied Energy, vol. 263, p. 114683, 2020.

[78] J. Vivian, E. Prataviera, N. Gastaldello, and A. Zarrella, "A comparison between grey-box models and neural networks for indoor air temperature prediction in buildings," Journal of Building Engineering, vol. 84, p. 108583, 2024.

[79] B. Amos, L. Xu, and J. Z. Kolter, "Input convex neural networks," in International Conference on Machine Learning. PMLR, 2017, pp. 146–155.

[80] S. Chen, K. Saulnier, N. Atanasov, D. D. Lee, V. Kumar, G. J. Pappas, and M. Morari, "Approximating explicit model predictive control using constrained neural networks," in 2018 Annual American Control Conference (ACC). IEEE, 2018, pp. 1520–1527.

[81] T. Hastie, R. Tibshirani, and J. Friedman, The elements of statistical learning. Springer, 2009.

[82] S. Geman, E. Bienenstock, and R. Doursat, "Neural networks and the bias/variance dilemma," Neural Computation, vol. 4, no. 1, pp. 1–58, 1992.

[83] X. Lin, Z. Tian, W. Song, Y. Lu, J. Niu, Q. Sun, and Y. Wang, "Grey-box modeling for thermal dynamics of buildings under the presence of unmeasured internal heat gains," Energy and Buildings, vol. 314, p. 114229, 2024.

[84] Z. Jiang and B. Dong, "Modularized neural network incorporating physical priors for future building energy modeling," Patterns, vol. 5, no. 8, 2024.

[85] Y. Li, Z. O'Neill, L. Zhang, J. Chen, P. Im, and J. DeGraw, "Grey-box modeling and application for building energy simulations — A critical review," Renewable and Sustainable Energy Reviews, vol. 146, p. 111174, 2021.

[86] Y. Liu, P. Kergus, F. Claveau, P. Chevrel, and B. Lacarrière, "Building thermal control: Hierarchical design from limited data using grey-box or black-box internal models for model predictive control," Journal of Building Engineering, p. 112936, 2025.

[87] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," Automatica, vol. 36, no. 6, pp. 789–814, 2000.

[88] A. S. Anand, S. Sawant, D. Reinhardt, and S. Gros, "Data-driven predictive control and MPC: Do we achieve optimality?" IFAC-PapersOnLine, vol. 58, no. 15, pp. 73–78, 2024.

[89] M. Gevers, A. S. Bazanella, X. Bombois, and L. Miskovic, "Identification and the information matrix: How to get just sufficiently rich?" IEEE Transactions on Automatic Control, vol. 54, no. 12, pp. 2828–2840, 2009.

[90] J. Drgoňa, A. R. Tuor, V. Chandan, and D. L. Vrabie, "Physics-constrained deep learning of multi-zone building thermal dynamics," Energy and Buildings, vol. 243, p. 110992, 2021.

[91] N. Asprion, R. Böttcher, R. Pack, M.-E. Stavrou, J. Höller, J. Schwientek, and M. Bortz, "Gray-box modeling for the optimization of chemical processes," Chemie Ingenieur Technik, vol. 91, no. 3, pp. 305–313, 2019.

[92] "Matlab Optimization Toolbox." [Online]. Available: https://fr.mathworks.com/help/optim/index.html

[93] M. Verhaegen and P. Dewilde, "Subspace model identification part 2. Analysis of the elementary output-error state-space model identification algorithm," International Journal of Control, vol. 56, no. 5, pp. 1211–1241, 1992.

[94] J. Cigler and S. Prívara, "Subspace identification and model predictive control for buildings," in 2010 11th International Conference on Control Automation Robotics & Vision. IEEE, 2010, pp. 750–755.

[95] G. Armenise, M. Vaccari, R. B. Di Capaci, and G. Pannocchia, "An open-source system identification package for multivariable processes," in 2018 UKACC 12th International Conference on Control (CONTROL). IEEE, 2018, pp. 152–157.

[96] E. F. Camacho and C. Bordons, "Constrained model predictive control," in Model predictive control. Springer, 2007, pp. 177–216.

[97] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," Mathematical Programming Computation, vol. 11, no. 1, pp. 1–36, 2019.

[98] V. Breschi and S. Formentin, "Direct data-driven control with embedded anti-windup compensation," in Learning for Dynamics and Control. PMLR, 2020, pp. 46–54.

[99] Y. Liu, P. Kergus, F. Claveau, and P. Chevrel, "Hierarchical Data-enabled Predictive Control for Building Thermal Management," in Control Applications of Optimisation (CAO), 2025.

[100] H. J. Ferreau, C. Kirches, A. Potschka, H. G. Bock, and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," Mathematical Programming Computation, vol. 6, pp. 327–363, 2014.

[101] J. C. Willems, "Paradigms and puzzles in the theory of dynamical systems," IEEE Transactions on Automatic Control, vol. 36, no. 3, pp. 259–294, 1991.

[102] I. Markovsky and F. Dörfler, "Behavioral systems theory in data-driven analysis, signal processing, and control," Annual Reviews in Control, vol. 52, pp. 42–64, 2021.

[103] C. De Persis and P. Tesi, "Formulas for data-driven control: Stabilization, optimality, and robustness," IEEE Transactions on Automatic Control, vol. 65, no. 3, pp. 909–924, 2019.

[104] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Data-driven model predictive control with stability and robustness guarantees," IEEE Transactions on Automatic Control, vol. 66, no. 4, pp. 1702–1717, 2020.

[105] H. J. Van Waarde, J. Eising, H. L. Trentelman, and M. K. Camlibel, "Data informativity: A new perspective on data-driven analysis and control," IEEE Transactions on Automatic Control, vol. 65, no. 11, pp. 4753–4768, 2020.

[106] M. Rotulo, C. De Persis, and P. Tesi, "Data-driven linear quadratic regulation via semidefinite programming," IFAC-PapersOnLine, vol. 53, no. 2, pp. 3995–4000, 2020.

[107] H. J. Van Waarde, M. K. Camlibel, and M. Mesbahi, "From noisy data to feedback controllers: Nonconservative design via a matrix S-lemma," IEEE Transactions on Automatic Control, vol. 67, no. 1, pp. 162–175, 2020.

[108] J. Berberich, C. W. Scherer, and F. Allgöwer, "Combining prior knowledge and data for robust controller design," IEEE Transactions on Automatic Control, vol. 68, no. 8, pp. 4618–4633, 2022.

[109] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Linear tracking MPC for nonlinear systems — Part II: The data-driven case," IEEE Transactions on Automatic Control, vol. 67, no. 9, pp. 4406–4421, 2022.

[110] C. Verhoek, R. Tóth, S. Haesaert, and A. Koch, "Fundamental lemma for data-driven analysis of linear parameter-varying systems," in 2021 60th IEEE Conference on Decision and Control (CDC). IEEE, 2021, pp. 5040–5046.

[111] T. Faulwasser, R. Ou, G. Pan, P. Schmitz, and K. Worthmann, "Behavioral theory for stochastic systems: A data-driven journey from Willems to Wiener and back again," Annual Reviews in Control, vol. 55, pp. 92–117, 2023.

[112] Y. Lian, J. Shi, M. Koch, and C. N. Jones, "Adaptive robust data-driven building control via bilevel reformulation: An experimental result," IEEE Transactions on Control Systems Technology, vol. 31, no. 6, pp. 2420–2436, 2023.

[113] J. Coulson, J. Lygeros, and F. Dörfler, "Regularized and distributionally robust data-enabled predictive control," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 2696–2701.

[114] C. Jeremy, J. Lygeros, and F. Dörfler, "Distributionally robust chance constrained data-enabled predictive control," IEEE Transactions on Automatic Control, vol. 67, no. 7, pp. 3289–3304, 2021.

[115] L. Huang, J. Coulson, J. Lygeros, and F. Dörfler, "Decentralized data-enabled predictive control for power system oscillation damping," IEEE Transactions on Control Systems Technology, vol. 30, no. 3, pp. 1065–1077, 2021.

[116] L. Huang, C. Jeremy, J. Lygeros, and F. Dörfler, "Data-enabled predictive control for grid-connected power converters," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 8130–8135.

[117] W. Favoreel, B. De Moor, and M. Gevers, "SPC: Subspace predictive control," IFAC Proceedings Volumes, vol. 32, no. 2, pp. 4004–4009, 1999.

[118] A. Chiuso, "The role of vector autoregressive modeling in predictor-based subspace identification," Automatica, vol. 43, no. 6, pp. 1034–1048, 2007.

[119] F. Fiedler and S. Lucia, "On the relationship between data-enabled predictive control and subspace predictive control," in 2021 European Control Conference (ECC). IEEE, 2021, pp. 222–229.

[120] F. Dörfler, J. Coulson, and I. Markovsky, "Bridging direct and indirect data-driven control formulations via regularizations and relaxations," IEEE Transactions on Automatic Control, vol. 68, no. 2, pp. 883–897, 2022.

[121] M. Yin, H. Cai, A. Gattiglio, F. Khayatian, R. S. Smith, and P. Heer, "Data-driven predictive control for demand side management: Theoretical and experimental results," Applied Energy, vol. 353, p. 122101, 2024.

[122] E. Elokda, J. Coulson, P. N. Beuchat, J. Lygeros, and F. Dörfler, "Data-enabled predictive control for quadcopters," International Journal of Robust and Nonlinear Control, vol. 31, no. 18, pp. 8916–8936, 2021.

[123] R. E. Kalman, "Design of a self-optimizing control system," Transactions of the American Society of Mechanical Engineers, vol. 80, no. 2, pp. 468–477, 1958.

[124] K. J. Åström, "Adaptive control," in Mathematical System Theory: The Influence of R. E. Kalman. Springer, 1995, pp. 437–450.

[125] J. Richalet, A. Rault, J. L. Testud, and J. Papon, "Model predictive heuristic control," Automatica, vol. 14, no. 5, pp. 413–428, 1978.

[126] C. R. Cutler, "Dynamic matrix control: an optimal multivariable control algorithm with constraints," PhD thesis, University of Houston, 1983.

[127] Clarke David W, and Mohtadi, Coorous and Tuffs, P Simon, "Generalized predictive control — Part I. The basic algorithm," Automatica, vol. 23, no. 2, pp. 137–148, 1987.

[128] Clarke David W, and Mohtadi, Coorous and Tuffs P Simon, "Generalized predictive control — Part II extensions and interpretations," Automatica, vol. 23, no. 2, pp. 149–160, 1987.

[129] M. Morari and J. H. Lee, "Model predictive control: past, present and future," Computers & Chemical Engineering, vol. 23, no. 4-5, pp. 667–682, 1999.

[130] J.-S. Kim, "Recent advances in adaptive MPC," ICCAS 2010, pp. 218–222, 2010.

[131] V. Adetola and M. Guay, "Robust adaptive MPC for constrained uncertain nonlinear systems," International Journal of Adaptive Control and Signal Processing, vol. 25, no. 2, pp. 155–167, 2011.

[132] M. Bujarbaruah, X. Zhang, M. Tanaskovic, and F. Borrelli, "Adaptive MPC under time varying uncertainty: Robust and Stochastic," arXiv preprint arXiv:1909.13473, 2019.

[133] W. Tang and P. Daoutidis, "Data-driven control: Overview and perspectives," in 2022 American Control Conference (ACC). IEEE, 2022, pp. 1048–1064.

[134] H. Chen and L. Guo, "Convergence of adaptive MPC for linear stochastic systems," Science China Information Sciences, vol. 66, no. 5, p. 152201, 2023.

[135] Y. Bao and J. M. Velni, "Adaptive Uncertainty Quantification for Scenario-based Control Using Meta-learning of Bayesian Neural Networks," IFAC-PapersOnLine, vol. 58, no. 28, pp. 486–491, 2024.

[136] R. Tao, P. Zhao, I. Kolmanovsky, and N. Hovakimyan, "Robust adaptive MPC using uncertainty compensation," in 2024 American Control Conference (ACC). IEEE, 2024, pp. 1873–1878.

[137] T. Salzmann, E. Kaufmann, J. Arrizabalaga, M. Pavone, D. Scaramuzza, and M. Ryll, "Real-time neural MPC: Deep learning model predictive control for quadrotors and agile robotic platforms," IEEE Robotics and Automation Letters, vol. 8, no. 4, pp. 2397–2404, 2023.

[138] S. J. Qin, "An overview of subspace identification," Computers & Chemical Engineering, vol. 30, no. 10-12, pp. 1502–1513, 2006.

[139] L. Sha, Z. Jiang, and H. Sun, "A control strategy of heating system based on adaptive model predictive control," Energy, vol. 273, p. 127192, 2023.

[140] C. Verhoek, J. Berberich, S. Haesaert, R. Tóth, and H. S. Abbas, "A linear parameter-varying approach to data predictive control," arXiv preprint arXiv:2311.07140, 2023.

[141] I. Houtzager, J.-W. van Wingerden, and M. Verhaegen, "Recursive predictor-based subspace identification with application to the real-time closed-loop tracking of flutter," IEEE Transactions on Control Systems Technology, vol. 20, no. 4, pp. 934–949, 2011.

[142] N. A. Mardi and L. Wang, "Subspace-based model predictive control of time-varying systems," in Proceedings of the 48th IEEE Conference on Decision and Control (CDC). IEEE, 2009, pp. 4005–4010.

[143] S. Baros, C.-Y. Chang, G. E. Colon-Reyes, and A. Bernstein, "Online data-enabled predictive control," Automatica, vol. 138, p. 109926, 2022.

[144] L. Schmitt, J. Beerwerth, M. Bahr, and D. Abel, "Data-driven predictive control with online adaption: Application to a fuel cell system," IEEE Transactions on Control Systems Technology, vol. 32, no. 1, pp. 61–72, 2023.

[145] P. Verheijen, G. R. G. da Silva, and M. Lazar, "Recursive data-driven predictive control with persistence of excitation conditions," in 2022 IEEE 61st Conference on Decision and Control (CDC). IEEE, 2022, pp. 467–473.

[146] Y. Wang and A. Singh, "An empirical comparison of sampling techniques for matrix column subset selection," in 2015 53rd Annual Allerton Conference on Communication, Control, and Computing. IEEE, 2015, pp. 1069–1074.

[147] J. Shi, Y. Lian, and C. N. Jones, "Efficient Recursive Data-enabled Predictive Control," in 2024 European Control Conference (ECC), 2024, pp. 880–887.

[148] K. Zhang, K. Chen, X. Lin, Y. Zheng, X. Yin, X. Hu, Z. Song, and Z. Li, "Data-enabled predictive control for fast charging of lithium-ion batteries with constraint handling," arXiv preprint arXiv:2209.12862, 2022.

[149] K. Zhang, Y. Zheng, C. Shang, and Z. Li, "Dimension reduction for efficient data-enabled predictive control," IEEE Control Systems Letters, vol. 7, pp. 3277–3282, 2023.

[150] A. Faye-Bedrin, S. Aranovskiy, P. Chauchat, and R. Bourdais, "A computationally efficient reformulation for Data-Enabled Predictive Control," in 2024 IEEE 63rd Conference on Decision and Control (CDC). IEEE, 2024, pp. 4855–4860.

[151] R. Ortega, J. G. Romero, and S. Aranovskiy, "A new least squares parameter estimator for nonlinear regression equations with relaxed excitation conditions and forgetting factor," Systems & Control Letters, vol. 169, p. 105377, 2022.

[152] L. Cao and H. Schwartz, "A directional forgetting algorithm based on the decomposition of the information matrix," Automatica, vol. 36, no. 11, pp. 1725–1731, 2000.

[153] R. Kulhavý, "Restricted exponential forgetting in real-time identification," Automatica, vol. 23, no. 5, pp. 589–600, 1987.

[154] A. L. Bruce, A. Goel, and D. S. Bernstein, "Recursive Least Squares with Matrix Forgetting," in 2020 American Control Conference (ACC), 2020, pp. 1406–1410.

[155] L. Z. Urrutia, J. F. Pascual, E. P. Iribarren, R. S. Garay, and I. G. Pino, "Model predictive control with self-learning capability for automated demand response in buildings," Applied Thermal Engineering, vol. 258, p. 124558, 2025.

[156] S. Sedghizadeh and S. Beheshti, "Data-driven subspace predictive control: Stability and horizon tuning," Journal of the Franklin Institute, vol. 355, no. 15, pp. 7509–7547, 2018.

[157] D. Ling, Y. Zheng, H. Zhang, W. Yang, and B. Tao, "Detection of model-plant mismatch in closed-loop control system," Journal of Process Control, vol. 57, pp. 66–79, 2017.

[158] A. Faye-Bédrin, P. Chauchat, S. Aranovskiy, and R. Bourdais, "Willems' Lemma Reformulations: Which Operators Preserve LTI System Behavior?" IEEE Control Systems Letters, 2025.

[159] M. S. Mirnaghi and F. Haghighat, "Fault detection and diagnosis of large-scale HVAC systems in buildings using data-driven methods: A comprehensive review," Energy and Buildings, vol. 229, p. 110492, 2020.

[160] K. Sunny, A. Sheikh, and S. Wagh, "Application of dynamic mode decomposition for temperature analysis in smart building," in 2020 7th International Conference on Control, Decision and Information Technologies (CoDIT), vol. 1. IEEE, 2020, pp. 1197–1202.

[161] X. Zhang, G. Schildbach, D. Sturzenegger, and M. Morari, "Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty," in 2013 European Control Conference (ECC). IEEE, 2013, pp. 1029–1034.

[162] W. Yu, Z. Tang, K. Pang, and J. Liu, "Centralized DeePC-Based Voltage Control for Networked Microgrids System," in 2023 2nd Asian Conference on Frontiers of Power and Energy (ACFPE). IEEE, 2023, pp. 351–355.

COLLEGES SCIENCES
BRETAGNE POUR L'INGENIEUR
LOIRE ET LE NUMERIQUE

IMT Atlantique
Bretagne-Pays de la Loire
École Mines-Télécom

Titre : Contrôle Prédictif Basé sur les Données pour la Gestion Thermique des Bâtiments : Un Cadre Hiérarchique Comparant les Approches Indirectes et Directes

Mots clés : Contrôle prédictif, Identification de systèmes, Contrôle hiérarchique, Contrôle Basé sur les données, Gestion thermique des bâtiments.

Résumé : Les bâtiments comptent parmi les plus grands consommateurs d'énergie, notamment pour le chauffage et la climatisation. La Commande Prédictive basée sur un Modèle (MPC) offre une voie prometteuse pour réduire cette consommation sans sacrifier le confort thermique, mais son déploiement à grande échelle se heurte à l'hétérogénéité des équipements et au manque de données opérationnelles. Cette thèse explore des approches prédictives basées sur les données pour relever ces défis. Un cadre de commande hiérarchique est proposé, où les contrôleurs de bas niveau sont ajustés via le Virtual Reference Feedback Tuning (VRFT), tandis qu'un modèle boîte grise est comparé à des modèles boîte noire (sous-espaces). Les essais montrent que la MPC boîte grise réalise des économies d'énergie substantielles même avec peu de données.

De plus, la Commande Prédictive Pilotée par les Données (DeePC) génère des prédictions directement à partir de trajectoires organisées en matrices de Hankel, sans identification explicite. Elle surpasse les contrôleurs superviseurs classiques, mais nécessite un réglage plus délicat et davantage de calculs. Enfin, pour faire face aux variations temporelles (saisons, vieillissement des équipements), des extensions adaptatives sont développées : réduction récursive par SVD rapide et fenêtre glissante de données. Ces méthodes améliorent significativement le compromis consommation-confort en conditions opérationnelles variables.

Title : Data-Driven Predictive Control for Building Thermal Management: A Hierarchical Framework Comparing Indirect and Direct Approaches

Keywords : Predictive control, System identification, Hierarchical control, Data-driven control, Building thermal management.

Abstract: Buildings account for a significant share of global energy consumption, especially for heating and cooling. Model Predictive Control (MPC) offers a promising way to reduce this consumption while maintaining thermal comfort, yet large-scale deployment faces two major challenges: equipment heterogeneity and limited operational data. This dissertation investigates data-driven predictive control approaches that extract system behavior from historical data to address these challenges. A hierarchical control framework is proposed, decoupling thermal planning from equipment regulation. Low-level controllers are tuned using Virtual Reference Feedback Tuning (VRFT) without explicit models. At the supervisory level, grey-box MPC is shown to achieve substantial energy savings even with limited data, outperforming black-box subspace identification methods.

Furthermore, Data-enabled Predictive Control (DeePC) generates predictions directly from input-output trajectories organized in Hankel matrices, requiring no explicit model identification. Unlike Subspace Predictive Control (SPC) with fixed offline models, DeePC constructs online dynamic representations via regularized trajectory combinations, outperforming supervisory controllers at the cost of more delicate tuning and higher computation. To handle time-varying dynamics caused by seasonal weather and equipment aging, adaptive extensions are developed—including recursive identification for SPC and sliding-window or forgetting-factor strategies for DeePC. Validation confirms these adaptive methods significantly improve the energy-comfort trade-off under varying operating conditions.