THÈSE

En vue de l'obtention du

DOCTORAT en SCIENCES

Spécialié : *Sciences et Technologies de l'Information et de la Communication*

Délivré par : *CY CERGY PARIS UNIVERSITÉ*

Présentée et soutenue le *05/02/2026* **par :**

MOUNCEF NAJI

Semantic Personalization of Informed Consent based on Knowledge Graph Reasoning and Large Language Models

---

JURY

MARIE-HÉLÈNE ABEL	Université de Technologie de Compiègne	Présidente
ZOLTAN MIKLOS	Université de Rennes	Rapporteur
MOURAD ABED	Université Polytechnique Hauts-de-France	Rapporteur
GAYO DIALLO	Université de Bordeaux	Examinateur
JEAN-MARIE JANUEL	Fondation Léonie Chaptal	Invité
HAJER BAAZAOUI	CY Cergy Paris Université	Directrice de thèse
MAROUA MASMOUDI	Université de Rennes	Co-Directrice de thèse

École doctorale et spécialité :

EM2PSI : Economie, Gestion, Mathématiques, Physique théorique et Informatique

Unité de Recherche :

Équipe Traitement de l'Information et Systèmes (ETIS)

Abstract

Informed consent - as a mutual agreement - is a cornerstone of patient autonomy and ethical practice in healthcare. Yet, despite legal foundations to ensure its necessity, current practices often are not capable of providing adequate basis for its' deployment. Furthermore, modern policies require a rethinking of informed consent approaches, in a way that satisfies it's legal and ethical aspects, integrated in a personalized healthcare pathway.

In this context, this thesis introduces SPIRIT — a Semantic and Personalized Compliant approach for informed Consent and care recommendation — a novel semantic-based approach that focuses on informed consent, starting from it's legal compliance, going through it's personalization, and finishing by care recommendation, all in favor of healthcare pathway. The proposed approach integrates ontological reasoning, legal compliance verification, and generative artificial intelligence to achieve three primary objectives: (1) facilitate legal and ethical consistency of consent documents through ontology-driven annotation and semantic similarity optimization; (2) adapt consent to patient profiles based on symbolic reasoning that incorporates knowledge graphs and large language models (LLMs); and (3) propose personalized nursing care plans through multi-label classification with clinical standards.

All aspects of the approach are evaluated by experiments that target document compliance, content personalization, and care recommendation. The results demonstrate how semantic integration and patient-centered reasoning improve document readability, flexibility, and consistency with medical and legal standards.

Finally, an online prototype that implements the full SPIRIT architecture is developed to render the approach functional and test it in actual healthcare practice. The prototype is a proof of concept for personalized and semantically compliant consent management that paves the way towards smart, ethical, and patient-centered digital health approaches.

Keywords: Informed Consent, Legal Compliance, Personalization, Knowledge Graph, Ontology, LLMs, NLP, Recommendation.

Résumé

Le consentement éclairé, en tant qu'accord mutuel entre patient et professionnel de santé, constitue l'un des fondements majeurs de l'autonomie du patient et de l'éthique clinique. Toutefois, malgré un cadre juridique solide qui en consacre l'exigence, les pratiques actuelles peinent encore à fournir les conditions nécessaires à une mise en œuvre réellement satisfaisante. Par ailleurs, l'évolution récente des politiques de santé impose de repenser les modalités mêmes du consentement éclairé, afin d'assurer conjointement le respect de ses dimensions légales et éthiques tout en l'intégrant pleinement dans un parcours de soins personnalisé.

Dans cette perspective, cette thèse propose SPIRIT — Semantic and Personalized Compliant approach for informed Consent and care recommendation — une approche innovante fondée sur des technologies sémantiques et dédiée au consentement éclairé. Elle couvre l'ensemble du cycle allant de la vérification de conformité juridique à la personnalisation du contenu, jusqu'à la recommandation de soins, dans une logique intégrée au parcours de santé. L'approche SPIRIT combine raisonnement ontologique, vérification de conformité légale et intelligence artificielle générative pour répondre à trois objectifs principaux : (1) garantir la cohérence juridique et éthique des documents de consentement grâce à une annotation ontologique et à une optimisation de similarité sémantique ; (2) adapter le contenu du consentement aux profils des patients au moyen d'un raisonnement sémantique articulant graphes de connaissances et grands modèles de langage (LLMs) ; (3) générer des plans de soins infirmiers personnalisés à l'aide d'une classification multi-labels fondée sur des standards cliniques reconnus.

L'approche a été évaluée à travers un ensemble d'expérimentations portant sur la conformité documentaire, la personnalisation du contenu et la recommandation de soins. Les résultats démontrent que l'intégration sémantique, associée à un raisonnement centré sur le patient, améliore significativement la lisibilité, la flexibilité et la cohérence des documents avec les exigences médicales et juridiques.

Enfin, un prototype en ligne implémentant l'architecture complète de SPIRIT a été développé afin de valider l'approche en situation réelle d'usage. Ce prototype constitue une preuve de concept robuste pour une gestion du consentement personnalisée et sémantiquement conforme, ouvrant la voie à de nouvelles pratiques numériques en santé, plus intelligentes, plus éthiques et véritablement centrées sur le patient.

Mots-clés : Consentement éclairé, Conformité légale, Personnalisation, Graphe de connaissances, Ontologie, LLM, NLP, Recommandation.

Remerciements

Accomplishing this thesis has been an exciting and demanding journey, marked by challenges, discoveries, and moments of profound learning. I would like to express my sincere appreciation to all those who contributed to this work and supported me throughout this remarkable adventure.

First and foremost, I extend my gratitude to all the members of the jury. I am especially thankful to Zoltan Miklos and Mourad Abed for kindly accepting to review this manuscript, and to Gayo Diallo and Sonia Yassa for their ongoing involvement and attentive follow-up throughout the course of my doctoral research.

I express my profound gratitude to Madame Hajer Baazaoui, my doctoral advisor, for giving me the opportunity to conduct this research under her supervision. I am particularly thankful for her continuous guidance, insightful feedback, and unwavering support. Her valuable suggestions and constructive comments have been a constant source of motivation and have greatly enriched this work.

My sincere thanks also go to Maroua Masmoudi, my co-supervisor, for her encouragement, expertise, and generosity. Her deep insights and advice have supported me at crucial stages of my research and beyond. I am extremely grateful for her ideas, thoughtful suggestions, stimulating discussions, and the inspiration she provided to always aim for improvement, both in work and life.

I would like to express my appreciation to Madame Maguelonne Legaie, General Director of the Fondation Léonie Chaptal, for welcoming me into the foundation, for her trust, and for allowing me to contribute to her vision and her work. Her valuable discussions helped shape and advance the direction of this thesis.

My thanks also go to Christophe Jaurand, Mourad Zghal, Guillaume Decormeille, and Jean-Marie Januel, whose contributions and exchanges at the foundation played a significant role in the development of this research.

I am sincerely grateful to all the members of the Fondation Léonie Chaptal for the numerous discussions, the questions they addressed, the feedback they provided, and, above all, the life lessons I had the privilege to learn from them.

My heartfelt thanks extend to all research scholars, colleagues, and friends for their support, and their helpful presence throughout this journey. I also wish to thank the administrative and technical staff of ETIS for their assistance and support, both academically and administratively, during these years of study.

Finally, this journey would not have been possible without the unconditional love and steady support of my family. I owe a profound debt of gratitude to my parents, my brothers and sisters, whose sacrifices, values, and constant support have shaped who I am and have carried me through every stage of this work. Their trust, their guidance, and their belief in my abilities have been my greatest source of strength. To them, I dedicate this achievement with my deepest respect and heartfelt appreciation.

With heartfelt gratitude,

Thank you.

Contents

Introduction **1**

1 Conceptual Background **7**

1.1 Introduction 8

1.2 Consent in Healthcare 8

1.2.1 Informed Consent: Concepts and Origins 8

1.2.1.1 Definitions and fundamental concepts 8

1.2.1.2 History of informed consent in Healthcare 10

1.2.1.3 Medical Informed Consent 11

1.2.1.4 Consent models 12

1.2.2 Challenges in Traditional Consent Processes 13

1.3 Background on AI 15

1.3.1 Evolution of AI 15

1.3.2 Generative AI and Language Models 16

1.3.2.1 Learning methods 16

1.3.2.2 Self-supervised pre-training approaches 17

1.3.2.3 Large Language Models 19

1.3.2.4 Training and Fine-tuning LLMs 22

1.3.2.5 Prompting LLMs 24

1.4 Knowledge Representation and Reasoning 29

1.4.1 Ontologies 30

1.4.1.1 Ontology Definitions 30

1.4.1.2 Ontology Construction 32

1.4.2 Semantic Web technologies and languages 33

1.4.2.1 Resource Description Framework (RDF) 34
1.4.2.2 Web Ontology Language (OWL) 35
1.4.2.3 Differences and Usage of Ontology Formats 35
1.4.2.4 SPARQL Protocol and RDF Query Language (SPARQL) . . 36
1.4.2.5 SWRL rules 37
1.4.2.6 Ontology programming tools 37
1.4.3 Knowledge graphs 37
1.5 Conclusion 38

2 SPIRIT : Semantic Personalized and complIant system for Informed consent and care RecommendaTion 41

2.1 Introduction 42
2.2 Semantic Personalized and complIant system for Informed consent and care RecommendaTion (SPIRIT) 42
2.2.1 Related works 42
2.2.2 SPIRIT architecture 45
2.2.2.1 Data acquisition and processing: 45
2.2.2.2 Personalized e-consent : 47
2.2.2.3 Medical recommendation : 48
2.2.3 Synthesis 49
2.3 Ontology-Driven machine learning annotation for Legal Compliance 49
2.3.1 Related works 50
2.3.1.1 NLP-based document analysis methods 50
2.3.1.2 Ontology-driven document representation 52
2.3.2 Ontology-Driven annotation architecture 54
2.3.2.1 Data Collection 55
2.3.2.2 Similarity and Annotation for Legal articles and consent Documents 56

2.3.3 Experimental results 59
2.3.3.1 Implementation Process 60
2.3.3.2 Experimental Setup 61
2.3.3.3 Parameter selection fo similarity algorithms 62
2.3.4 Knowledge graph embedding and annotation 62
2.3.5 Synthesis 64
2.4 Conclusion 65

3 Semantic Personalization of Informed Consent 67

3.1 Introduction 67
3.2 Related Works 69
3.2.1 LLM-Based Personalization Systems 69
3.2.2 Profile ontologies 71
3.2.3 Identified Limitations and Contribution of Our Work 75
3.3 Semantic Personalization over LLMs 75
3.3.1 Formal problem definition 76
3.3.2 Retrieval Layer 77
3.3.3 Semantic Paths Extraction Layer 78
3.3.4 Semantic Reasoning Layer 81
3.3.5 Dynamic Prompt Engineering Layer 84
3.4 Experimental Results 85
3.4.1 Datasets 86
3.4.2 Environment configuration 86
3.4.3 Evaluation metrics 87
3.4.3.1 Statistical Approaches 87
3.4.3.2 Model-Based (Non-LLM) Approaches 88
3.4.3.3 LLM-Evaluation Approaches 88

3.4.4 Evaluation Setup . . . 89
3.4.5 Results . . . 92
3.5 Conclusion . . . 94

4 Care Recommendation **95**

4.1 Introduction . . . 95
4.2 Ontology-Driven Care Recommendation . . . 96
4.2.1 Related Works . . . 97
4.2.1.1 Medical Knowledge Graphs . . . 97
4.2.1.2 User Profiles in Personalized Healthcare . . . 98
4.2.1.3 Multi-Label Classification for Complex Medical Decision-Making 98
4.2.1.4 Synthesis . . . 99
4.2.2 Care recommendation and planning system . . . 99
4.2.2.1 Input Layer . . . 101
4.2.2.2 Parameter Augmentation . . . 102
4.2.2.3 Building the Knowledge Graph from Local Data . . . 103
4.2.2.4 Multi-label classification . . . 105
4.2.2.5 Classification model . . . 106
4.2.2.6 Interpretation to patient . . . 106
4.3 Experiments and Results . . . 106
4.3.1 Knowledge Graph Construction and Analysis . . . 107
4.3.2 Multi-Label Classification Performance . . . 108
4.3.3 Use case: Personalized care plan . . . 109
4.4 Conclusion . . . 110

5 Conclusions et perspectives **111**

5.1 Summary . . . 111

5.2 Limitations 113
5.3 Perspectives 114

A Proof of Concept **115**

A.0.1 System Overview and Architecture 115
A.0.2 Personalization of Consent Documents 116
A.0.3 Care Planning and Calendar Module 117
A.0.4 Security and Perspectives 117
A.0.5 Conclusion 118

Bibliographie **129**

List of Figures

1 Mind Map of our research methodology 5

1.1 Tim Berners-Lee's Semantic Web Stack [68]. 34

1.2 Graph representation of an RDF triple. 35

2.1 SPIRIT : Semantic Personalized and Compliant System for Informed Consent and care Recommendation 46

2.2 Ontology-Driven machine learning annotation for Legal Compliance 56

2.3 Overview of the LOINC Document Ontology axes it's concepts. 58

2.4 Parameter selection for similarity algorithms 62

(c) Cosine similarity vs. max length for XLNet 62

(d) Cosine similarity vs. window size for Word2Vec and FastText 62

2.5 Similarity measures for methods with and without Knowledge graph 63

2.6 Generated similarity annotations 64

(a) XLNET results 64

(b) Similarity measures 64

3.1 Semantic Personalization over LLMs approach 76

3.2 Concepts of CSTDO visualized in Protégé 80

3.3 Performance comparison of text personalization without and with ontology guidance. 93

3.4 LLM-as-a-judge on reasoning performances 94

4.1 Care Recommendation and planning system. 102

4.2 Performance metrics across 10-fold cross-validation: (Left) F1 Micro and AUC Micro; (Right) Hamming Loss 108

4.3 Use case: Personalized care plan 109

A.1 Dashboard of the application 116
A.2 Document personalization 117
A.3 Care planning 118

List of Tables

1.1 In each step, the model selects the token $x_i \in V$ that maximizes the conditional probability given the current context. 20

2.1 Related works on informed consent solutions 44

2.2 Comparative Analysis of Text Similarity Methods 53

2.3 Legifrance API Response Metadata 57

3.1 Comparison of Knowledge Graph-Augmented Retrieval Techniques using LLMs. 74

3.2 Consistency Scoring Criteria 90

3.3 Faithfulness Scoring Criteria 91

3.4 Coverage Scoring Criteria 91

3.5 Clarity / Comprehensibility Scoring Criteria 91

4.1 Summary of state-of-the-art investigations on data representation, ontology usage, transformers, classification types, and results achieved in medical decision systems 100

4.2 Knowledge Graph Evaluation 107

4.3 Multi-Label Classification Results 108

List of Acronyms

ASAD	Aide et Soins à Domicile (Home Care and Assistance Service)
SSIAD	Service de Soins Infirmiers À Domicile (Home Nursing Care Service)
EHR	Electronic Health Record
ICNP	International Classification for Nursing Practice
LOINC	Logical Observation Identifiers Names and Codes
ICO	Informed Consent Ontology
GDPR	General Data Protection Regulation
SW	Semantic Web
KG	Knowledge Graph
RDF	Resource Description Framework
OWL	Web Ontology Language
SWRL	Semantic Web Rule Language
SPARQL	SPARQL Protocol and RDF Query Language
JSON	JavaScript Object Notation
RDFLib	RDF Library for Python
AI	Artificial Intelligence
ML	Machine Learning
NLP	Natural Language Processing
LLM	Large Language Model
GPT	Generative Pre-trained Transformer
RAG	Retrieval-Augmented Generation
API	Application Programming Interface
REST	Representational State Transfer
CSV	Comma-Separated Values
PDF	Portable Document Format

Introduction

Context and motivations

Healthcare systems are nowadays bound to transformations, led by advances in medical technologies, treatments complexity, and the increasing shift towards patient-centered care. In this prospect, the relationship between healthcare providers and patients has shifted from a paternalistic model—where decisions were largely dictated by clinicians—to one where patients are recognized as active participants in their own care. This transition reflects broader values that emphasise autonomy, transparency, and shared decision-making. At the heart of this transformation lies the process of *informed consent*, which ensures that medical interventions are not only clinically appropriate but also ethically and legally grounded in the patient's understanding and voluntary agreement.

Informed consent is a fundamental aspect of modern medicine and biomedical ethics. It expresses the principles of autonomy, beneficence, and justice that have guided medical ethics since the *Belmont Report* (1979) and the *Declaration of Helsinki* (World Medical Association, 2013). The basic idea is that patients have the right to make decisions about their care treatment based on adequate information, free from coercion, and with a clear understanding of risks, benefits, and alternatives [1, 2].
Legally, informed consent has become embedded in most healthcare systems worldwide. In France, for example, the 2002 *"Loi relative aux droits des malades et à la qualité du système de santé"* (often called the loi Kouchner) expresses that "no medical act may be carried out without the free and informed consent of the person, and this consent may be withdrawn at any time." Similar obligations exist across Europe, the U.S., and internationally, implying informed consent as both a legal duty and a fundamental patient right [3], which guaranties patients' protection, autonomy, and builds trust in the clinician–patient relationship.

Despite its importance, there is a significant gap between the ideal model of informed consent and its actual practice. Indeed, empirical studies show that often patients do not fully understand the provided information during the consent process, and do not contribute to their care pathway. For instance, participants' comprehension in clinical trials has been reported to range from roughly 52% to 76% depending on the study and the component measured [4, 5]. This is reflected in many contributing factors:

- Readability barriers: Consent documents are written well above the average reading level of the population, with no adaptation to different literacy levels.

- Length and complexity: Consent documents often contain many pages, filled with legal terms and dense medical terminology. This compromises comprehension, especially in situations where patients are asked to sign quickly.

- Vulnerable groups: Elderly, patients with cognitive impairments (dementia, Alzheimer), or those with limited literacy and language barriers are disproportionately affected.

As a result, consent in practice is often reduced to a formal signature process rather than a genuine process of understanding.

In the meanwhile, on one hand, standard terminologies and knowledge graphs showcase the rapid digital transformation that is in place to empower interoperability with a common ground for both legal and medical practices, informed consent included. The *Informed Consent Ontology (ICO)* [6] and *SNOMED CT* [7] are some of the building blocks for such a unified model. Yet, the adaptive and personalization aspects of the consent process still present an issue in literature.

On the other hand, recent advances in natural language processing through *Large Language Models (LLMs)* such as GPT, LLaMA, or PaLM, have demonstrated remarkable capabilities in content understanding and generation [8, 9]. Indeed, technologies such as *retrieval-augmented generation (RAG)* [10] and *chain-of-thought prompting (CoT)* [11], present new aspects in handling the previous issues. Yet, LLMs alone remain limited in reasoning, explainability, and compliance assurance—crucial factors for legal and ethical contexts such as informed consent, due to the lack of grounding in structured, domain-specific knowledge.

In short, there is a clear and pressing need to transform consent from a static, legalistic document into an adaptive, personalized, and patient-centered interaction. Existing studies document limited comprehension and low readability of consent materials; at the same time, advances in medical terminologies and in AI provide the technical ingredients for a new generation of consent systems. Yet, literature lacks a unified approach that integrates the legal, ethical, and clinical dimensions of informed consent within a single approach — a gap this thesis aims to address.

CIREX - Fondation Léonie Chaptal

This doctoral study has been conducted as part of the CIREX research chair (Chaire Intelligence, Recherche et eXploration), in contractual partnership with the "Fondation Léonie Chaptal". The CIREX chair [12] is aimed at reconciling artificial intelligence research and applied healthcare innovation, by creating links between academic research and healthcare organizations. Its goal is to study the ethical, legal, and technological dimensions of AI in health, with special attention to transparency, explainability, and conformity in intelligent decision-support systems.

The Fondation Léonie Chaptal [13], established in 1913, is a medico-social nonprofit organization renowned for its commitment to professional training as well as home care services. The organization consists of two main branches: a training branch, encompassing nursing and healthcare training centers (IFSI, IFAS, IFAP), and a care branch, providing home-based medical and social assistance through services such as SSIAD (Home Nursing Care Service)

and ESAD (Specialized Home Care Service). Its practice is guided by a deep humanist and ethical philosophy, with the focus on patient dignity, autonomy, and the quality of care.

Within this collaboration, the Fondation Léonie Chaptal provided the applied field context for this study, with the focus on digitalization and personalization of informed consent processes in home care. The collaboration aimed to solve the day-to-day problems of medical professionals in collecting, managing, and customizing consent information to patients who possess diverse medical, cognitive, and social profiles. This scenario was the origin of both the data model and the field experiment site for the proposed semantic and AI-based solutions developed under this thesis.

Thesis Objectives

In this thesis, we focus on the challenges of personalized informed consent, we tackle the following question :

" How can we design a system capable of producing a medical consent that is simultaneously legally compliant, comprehensible, and adapted to the patient's profile, while supporting healthcare professionals in proposing appropriate care plans? "

Addressing this question requires tackling multiple disciplines

- **Legal compliance:** ensuring that the generated consent content conforms to applicable ethical and legal frameworks (e.g., the French Loi Kouchner, GDPR, and international biomedical standards);
- **Personalized consent and patient-centered:** adapting the consent information to the patient's profile characteristics, such as age, literacy level, pathology, and language;
- **Medical assistance:** allowing healthcare professionals to use consent information as an entry point for care pathway recommendations.

This work therefore seeks to move beyond the traditional static model of informed consent by introducing a hybrid AI-based approach—leveraging symbolic reasoning through Knowledge Graphs (KGs) and generative adaptability through Large Language Models (LLMs)—to create an *intelligent consent ecosystem.*

Ultimately, the goal is to enable a dynamic consent generation process that evolves with the patient's context, supports comprehension and autonomy, and ensures both ethical and legal compliance.

Research Questions

Building on the thesis objectives, this work is structured around three core research questions, each corresponding to one of the main components of the proposed approach:

1. **RQ1 – Legal compliance:** How can we ensure the automatic verification of consent documents against applicable legal and ethical standards?
2. **RQ2 – Personalization:** How can informed consent documents be adapted to individual patient profiles through semantic and medical knowledge?
3. **RQ3 – Care recommendation:** How can patient data and clinical ontologies be leveraged to recommend personalized care plans?

Contributions

The objective of this thesis lies in the conception of a **novel approach of informed consent** that integrates legal, semantic, and medical recommendation into a unified, interoperable, and patient-centered framework. Unlike traditional models that treat informed consent as a static and administrative formality, our approach presents it as a dynamic, intelligent, and adaptive process—one that ensures legal compliance, personalizes content according to patient profiles, and supports care professionals through data-driven recommendations.

To put into action the proposed approach, the thesis proposes three interconnected technical contributions, each addressing one of the research questions introduced earlier.

- **C1 – Semantic similarity enhancement for legal conformity:** We propose an ontology-based method for verifying the legal compliance of consent documents by computing semantic similarity between document paragraphs and legal texts. The approach leverages knowledge graph representations and hyperparameter tuning of similarity algorithms to enhance the semantic alignment between regulatory articles (e.g., from Légifrance) and consent documents. This ensures that the generated or used documents remain consistent with legal and ethical standards, providing an automated, explainable layer of compliance control.
- **C2 – Semantic prompting for consent personalization:** We introduce a hybrid reasoning–generation pipeline that combines symbolic reasoning from knowledge graphs with the generative capabilities of large language models (LLMs). Through SWRL rules and ontology reasoning, patient-specific personalization factors are inferred and dynamically integrated into LLM prompts. This contribution establishes a new paradigm of *semantic prompting* for personalized consent generation, grounded in structured domain knowledge and ethical awareness.

- **C3 – Nursing recommendation through multi-label classification:** We develop a predictive model for generating personalized nursing care recommendations based on patient profiles and standardized ontologies such as ICNP and SNOMED CT. By aligning local institutional data with international standards, the model not only improves predictive accuracy but also ensures interoperability and knowledge transfer between healthcare systems. The resulting classification system assists healthcare professionals in identifying relevant care actions tailored to each patient's needs.

Finally, as a proof of concept, this thesis proposes a **web-based prototype implementation** of the global approach. The application can be directly used by healthcare professionals. It provides an interface for testing the system's functionalities, regarding consent personalization and care plan suggestion. This prototype serves as both a validation and evaluation framework, enabling field experiments and user feedback collection within real healthcare environments. As well as its usability, interpretability, and acceptance among healthcare providers.

Together, these contributions advance the state of the art by bridging three traditionally isolated dimensions—*legal compliance*, *content personalization*, and *clinical recommendation*—into a coherent, ontology-driven framework. This unified approach lays the foundation for a new generation of informed consent systems that are **legally sound, ethically aligned, scientifically personalized, and patient-centered.** Figure 1 illustrates the mind map of the research methodology.

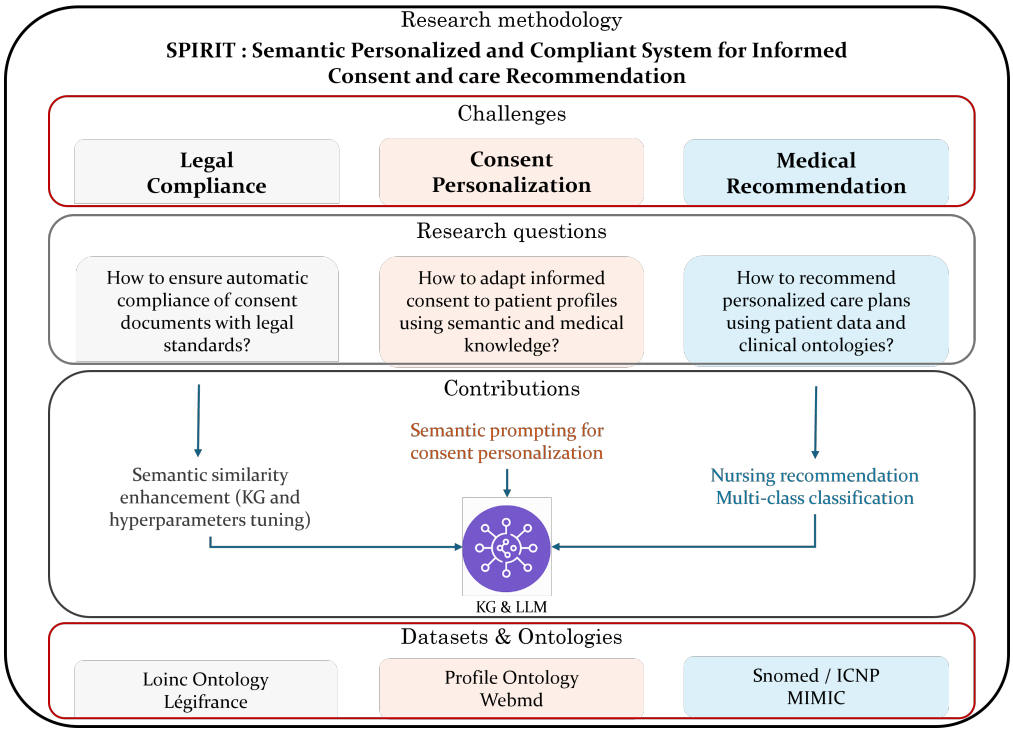


Figure 1 – Mind Map of our research methodology

Thesis outline

Chapter 1 presents the theoretical and conceptual foundations underlying the research. It provides an overview of the key domains involved, including knowledge graphs and ontologies,large language models, and informed consent in healthcare. The objective is to establish the necessary background for understanding the subsequent methodological and technical developments.

Chapter 2 is divided into two parts. The first part introduces "SPIRIT" - the global architecture of the proposed system - by presenting an overview of related works regarding consent systems, followed by a detailed description of "SPIRIT". In the second part, we delve into the first contribution of our work regarding legal compliance. We first give a state of the art of semantic similarity techniques, then we present our approach "Ontology-Driven machine learning annotation for Legal Compliance". The approach is then evaluated through different experiments based on cosine similarity calculations for different setups.

Chapter 3 focuses on the second contribution of our proposal, which addresses the personalization of consent documents according to patient profiles. First, we give an overview of related work research on LLMs and personalization. Then we present our approach "SPeLL" - Semantic Personalization over LLMs - that describes how we tackle the issue of personalizing consent by using ontologies over LLMs. Finally, we evaluate the approach using several state of the art metrics, as well as different perspectives, essentially with and without the use of ontologies.

Chapter 4 presents the third contribution of this thesis, which focuses on personalized care recommendation. We first introduce the state of the art in recommendation systems and multi-label classification in the medical domain. We then describe our proposed approach, which integrates clinical ontologies (ICNP and SNOMED-CT), a knowledge graph constructed from institutional nursing standards, and a transformer-based multi-label classifier enriched with ontology-based features. We further explain how the knowledge graph is built, how embeddings are derived, and how predictions are interpreted using LLM-based explanation. Finally, we report evaluation results for both the knowledge graph and the care recommendation engine.

Finally, in **Chapter 5** , we provide concluding remarks, the limitations of our work, and potential perspectives for future work. An implementation of the proposed workflow, including personalized consent generation and care planning, is provided as a Proof of Concept in Appendix A. The appendix contains the full system architecture, web-based interface, and integration details demonstrating the practical application of the approach in a home-care setting.

Chapter 1

Conceptual Background

Sommaire

1.1 Introduction 8
1.2 Consent in Healthcare 8
1.2.1 Informed Consent: Concepts and Origins 8
1.2.1.1 Definitions and fundamental concepts 8
1.2.1.2 History of informed consent in Healthcare 10
1.2.1.3 Medical Informed Consent 11
1.2.1.4 Consent models 12
1.2.2 Challenges in Traditional Consent Processes 13
1.3 Background on AI 15
1.3.1 Evolution of AI 15
1.3.2 Generative AI and Language Models 16
1.3.2.1 Learning methods 16
1.3.2.2 Self-supervised pre-training approaches 17
1.3.2.3 Large Language Models 19
1.3.2.4 Training and Fine-tuning LLMs 22
1.3.2.5 Prompting LLMs 24
1.4 Knowledge Representation and Reasoning 29
1.4.1 Ontologies 30
1.4.1.1 Ontology Definitions 30
1.4.1.2 Ontology Construction 32
1.4.2 Semantic Web technologies and languages 33
1.4.2.1 Resource Description Framework (RDF) 34
1.4.2.2 Web Ontology Language (OWL) 35
1.4.2.3 Differences and Usage of Ontology Formats 35
1.4.2.4 SPARQL Protocol and RDF Query Language (SPARQL) 36
1.4.2.5 SWRL rules 37
1.4.2.6 Ontology programming tools 37
1.4.3 Knowledge graphs 37
1.5 Conclusion 38

1.1 Introduction

Healthcare decision-making relies heavily on the ability to communicate medical information in a clear, trustworthy, and patient-centered manner. This communication lean on the concept of informed consent, a legal and ethical obligation that ensures patients understand and voluntarily agree to medical procedures and the use of their personal data. However, despite its fundamental role, traditional consent processes often remain generic, difficult to understand, and poorly adapted to individual patient needs and contexts. This creates a gap between regulatory requirements and real clinical practice. In parallel, healthcare has witnessed the rapid evolution of Artificial Intelligence (AI), particularly with the emergence of Large Language Models (LLMs) and knowledge-based systems. These technologies offer new opportunities to enhance patient information, personalize consent documents, and support clinical decision-making. Yet, their integration must be grounded in a solid understanding of the legal, ethical, and clinical foundations of consent, as well as reliable methods for representing medical knowledge.

This chapter provides the conceptual background necessary to understand the scientific, legal, and technological context of personalized consent in healthcare. First, it examines the origins and principles of informed consent, followed by its legal, ethical, and clinical implications, and the limitations of current practices (Section 1.2). Then, it explores the evolution of AI in healthcare, with a focus on generative AI and language models (Section 1.3). Finally, it presents knowledge representation approaches—such as ontologies and knowledge graphs, as well as semantic web technologies and querying languages in (Section 1.4).

1.2 Consent in Healthcare

The question of free and informed consent is nowadays a fundamental question regarding ethics and dignity of people. In this section, we cover the different concepts related to consent. First, we define what a consent is in Section 1.2.1.1. Then, we trace the history of informed consent across time with their application domains in Section 1.2.1.2. Afterwards, we present the different consent models in Section 1.2.1.4. Finally, we detail the challenges of getting and applying informed consent in practice in section 1.2.2.

1.2.1 Informed Consent: Concepts and Origins

1.2.1.1 Definitions and fundamental concepts

This subsection examines how the concepts of free and informed consent are defined and interpreted across various domains, including academic literature, legal regulations, and bioethics.

a. Dictionary perspective

To better understand the concept of informed consent, it is useful to examine the meaning of its core components as defined in standard dictionaries:

- **Consent** is defined as the "act of agreeing to an action or a plan."
- **Free** means "not subject to any constraint, control, or restriction."
- **To inform** means "to make something clearer, more understandable, or to explain it."

These definitions highlight that informed consent is not only about agreeing, but about agreeing freely, after having received and understood relevant information.

b. Legal perspective

The concept of medical consent entered French legal discourse with Law n°2002-303 of March 4, 2002, known as the *"Kouchner Law"*[1], which established the rights of patients and reinforced the quality of the healthcare system. This law represented a significant change from a paternalistic approach—where doctors made decisions for patients—toward a system that emphasizes the patient's autonomy and active participation.

Article L1111-4 of the French Public Health Code (CSP) states that:

> *"Every person makes decisions regarding their health together with the healthcare professional, considering the information and recommendations provided."*

This article confirms that the patient has the right to accept or refuse any medical intervention, and that healthcare professionals are required not only to respect these choices but also to adapt medical care accordingly. Furthermore, the patient must be clearly informed of the risks, consequences, and potential severity of their decisions.

Moreover, the same article emphasizes:

> *"No medical act or treatment can be carried out without the free and informed consent of the person, and this consent can be withdrawn at any time."*

This underlines two fundamental principles: consent must be both free (given without coercion or pressure) and informed (based on clear and understandable information). In addition, consent is not fixed—it is a dynamic and reversible process. A patient who initially agrees

[1] Kouchner law : "LOI n° 2002-303 du 4 mars 2002 relative aux droits des malades et à la qualité du système de santé (1)"

to a treatment retains the right to refuse or withdraw consent at a later stage. In summary, the Kouchner Law legally establishes medical consent as a continuous, autonomous decision-making process, grounded in respect for the patient's dignity, freedom of choice, and right to be informed.

c. Bioethical perspective

Bioethics is the study of ethical, social, and legal issues that arise in biomedicine and biomedical research (NIHS). Defining consent in bioethics is essential since it ensures patient autonomy, provides clear information for decision-making, and guides ethical practice in clinical and research settings. In that context, two forms of informed consent are identified [1]:

- **Active consent:** This form emphasizes the patient's autonomy and ability to make informed choices. The patient is actively involved and voluntarily agrees after understanding all relevant information (CIOMS [2])
- **Passive or presumed consent:** Here, consent is assumed in the absence of a refusal. It is not considered a genuine expression of agreement. This is applied in situations where a person is no longer able to express their will, such as brain death.

1.2.1.2 History of informed consent in Healthcare

The concept of consent in healthcare has evolved over centuries through a convergence of ethical, legal, and medical developments. One of the earliest known texts addressing medical ethics is the Hippocratic Oath, written in the 4^{th} century BCE. It served both as a moral contract for students of medicine and as a code of professional conduct for physicians, emphasizing above all the duty to "do no harm." Other foundational texts followed. Maimonides' Prayer, for instance, highlighted the inseparability of the body and mind, advocating for a holistic approach to healing [14].

In the modern era, the French Declaration of the Rights of Man and of the Citizen (1789) laid a legal foundation for ethical medical practice, emphasised by the Article 4 of the declaration. Key judicial decisions in the 20^{th} century have been taken for informed consent. The Mercier ruling of 1936 recognized the contractual nature of the doctor–patient relationship. This decision introduced the idea that doctors must inform patients about their care options. The 1942 Teyssier ruling required physicians to explain the benefits, drawbacks, and risks of various treatment options.

The explicit legal recognition of consent emerged after the World War II with the 1947 Nuremberg Code. This document created during the trials of Nazi physicians accused of war

[2]CIOMS : Council for International Organizations of Medical Sciences.

crimes, asserted that "the voluntary consent of the human subject is absolutely essential." It emphasized that consent must be given freely, without coercion or deception, and with full knowledge of the nature and implications of the procedure or experiment. This was followed by the 1948 Declaration of Geneva by the World Medical Association, the 1948 Universal Declaration of Human Rights, and the 1949 International Code of Medical Ethics, covering the same principles.

The 1964 Declaration of Helsinki introduced the concept of vulnerable populations in the context of medical research. Article 19 defines vulnerable individuals as those with a heightened risk of abuse or harm and mandates appropriate protections.

In France, legal statutes have started treating these ethical principles. The 1988 Huriet-Sérusclat Law was the first French bioethics law, aimed at protecting individuals participating in medical research. It clearly established the healthcare provider's obligation to inform and to obtain consent. The 2002 Kouchner Law advanced patients' rights by granting them full access to their medical records and emphasizing transparency in care. Finally, the 2005 Léonetti Law addressed end-of-life care, affirming that a patient's wish to limit or stop treatment must be respected and outlawing disproportionate therapeutic obstinacy.

1.2.1.3 Medical Informed Consent

Empirical research on medical informed consent consistently highlights persistent gaps between ethical requirements and clinical practice across healthcare systems worldwide. Systematic evidence indicates that many patients fail to fully understand essential elements of informed consent, including risks, benefits, alternatives, and uncertainty. A large systematic review by Pietrzykowski and Smilowska demonstrated that despite formal consent procedures, patient comprehension remains limited in many clinical and research settings, raising concerns about the validity of consent as an expression of autonomous decision-making [15]. Earlier work by Flory and Emanuel similarly showed that standard consent documents alone are insufficient, while targeted communication interventions may moderately improve understanding, supporting a shift toward more interactive and patient-centered consent processes [16].

International consensus-building efforts have increasingly focused on defining measurable outcomes for informed consent quality. Convie et al. conducted a multi-country stakeholder study to develop a core outcome set for informed consent in therapeutic contexts, emphasizing outcomes such as understanding, voluntariness, and alignment with patient values [17]. These findings resonate with shared decision-making frameworks promoted in clinical research, including those advanced by Légaré and colleagues, where informed consent is viewed not as a discrete legal event but as a communicative process integrated into care.

Empirical studies have also examined clinician-related factors influencing consent quality. A cross-sectional study in Romania revealed substantial variability in physicians' legal knowledge regarding informed consent and confidentiality, suggesting that institutional and educational factors play a critical role in ensuring ethically compliant practice [18]. Similarly,

observational research conducted in Spain assessed the formal quality of informed consent documents in critical care and surgical settings and found frequent deficiencies in personalization, risk disclosure, and documentation completeness, indicating systemic shortcomings in consent implementation [19].

Specific patient populations introduce additional complexity into informed consent practices. Research in pediatric care contexts, such as a recent study conducted in Cameroon, highlighted ethical and legal tensions in obtaining consent through parents or guardians, particularly in environments characterized by limited resources and variable health literacy [20]. These findings underscore the need for adaptive consent models sensitive to vulnerability, cultural context, and decisional capacity.

Within the French and broader European context, empirical studies have documented similar challenges. Dunn et al. identified therapeutic misconception among participants in French clinical research, demonstrating that participants often conflate research objectives with individualized clinical care [21]. Complementing this work, Hénin et al. proposed structured guidelines for drafting written information and consent documents in biomedical research, aiming to improve clarity, legal compliance, and participant understanding [22]. More recently, Glauzy et al. argued that the rise of precision medicine further complicates informed consent by increasing informational complexity, reinforcing the need for iterative, dialogical consent processes rather than static documentation [23].

Collectively, these international empirical studies support a reconceptualization of informed consent as a dynamic, context-sensitive, and patient-centered process. This perspective aligns closely with shared decision-making approaches and highlights the relevance of developing adaptive, personalized, and evaluable consent mechanisms in modern healthcare systems.

1.2.1.4 Consent models

Informed consent reflects the relationship between the doctor and the patient [24, 25]. Different models have been proposed over time, each one reflecting a different understanding of roles and responsibilities. In this section, we present four major models of consent.

The Paternalistic Model Also known as the parental or priestly model, the paternalistic approach centers on the doctor's responsibility to promote the patient's health and well-being. The doctor, based on clinical expertise, identifies the optimal intervention and presents only the information that supports this choice, encouraging the patient's consent.

The Informative Model Also referred to as the scientific, engineering, or consumer model, this approach focuses on giving the patient with all relevant medical facts. The doctor provides a detailed description of the diagnosis, treatment options, associated risks

and benefits, and uncertainties, enabling the patient to make fully informed decisions aligned with their personal values.

The Interpretive Model In the interpretive model, the doctor's role is not limited to presenting facts; it also involves helping the patient understand and clarify their values in the context of their medical condition. The doctor collaborates with the patient to explore goals, aspirations, and priorities, aiming to propose a treatment aligned with the patient's values.

The Deliberative Model The deliberative model represents the interaction between the doctor and the patient as a dialogue about the most appropriate health-related values and how to realize them through medical decisions. The goal is to engage in moral deliberation regarding which values and treatments best promote the patient's well-being, without coercion.

While each model aims to preserve patient autonomy, they differ in how autonomy is conceptualized—ranging from passive assent (paternalistic) to active moral self-development (deliberative). No single model is universally superior; instead, their relevance may vary depending on the clinical context and the patient's needs and capacities. Currently, at the Léonie Chaptal Foundation, it is observed that, in practice, the Paternalistic model is often followed for the acquisition of informed consent. Even though the interpretive model (collaborative) is the model that the foundation aims for, it is observed that it is hard to put it in place. This difficulty arises from the challenges of language, limited time and the complexity of the documents presented to patients. One of the purposes of this thesis is to help the transition from the paternalistic model to an interpretive one, by addressing the mentioned challenges.

1.2.2 Challenges in Traditional Consent Processes

The challenge of acquiring informed consent is linked to the patient's incapacity of consenting, due to physical, mental, or social factors, hindering their ability to make autonomous decisions. According to the World Health Organization (WHO) [3], vulnerability implies a partial or complete inability to protect oneself, often influenced by factors such as decision-making power, education level, access to resources, or physical strength.

In addition to inherent patient vulnerabilities, practical challenges can hinder the consent process. Consent forms are often lengthy and written in technical language, which can be difficult to understand. Limited time with healthcare professionals, patient stress or fatigue, and the cognitive load of processing complex medical information further reduce the ability to make an informed decision.

Communication barriers can also affect consent. Differences in language, cultural background, or health literacy may prevent patients from fully understanding the implications of

[3] World Health Organization

medical interventions.

Despite these practical challenges, the legal system generally presumes that individuals are capable of making decisions. From a legal standpoint, a person is presumed capable unless proven otherwise. According to the French Civil Code (Articles 1145 - 1146), all individuals can legally engage in contracts unless declared incapable. It specifies that unemancipated minors and adults under protection are considered legally incapable.

In practice, capacity is assumed unless the individual exhibits behavior that suggests otherwise—particularly in medical decision-making. However, there is no precise legal definition of how capacity should be assessed, which can result in inconsistent practices. Declaring someone incapable serves as a protective mechanism, not a punitive one.

The concepts of capacity and incapacity help define a person's level of autonomy. Without clear evaluation criteria, healthcare systems risk reverting to paternalistic models, where professionals override patients' wishes in the name of protection.

[1] proposed a framework in 2008 for evaluating decision-making capacity. Their model outlines a list of incapacity, ranked from least to most serious:

1. Inability to express or communicate a preference.
2. Inability to understand one's situation and its implications.
3. Inability to comprehend relevant information.
4. Inability to provide a reason for a choice.
5. Inability to provide a rational reason, even when secondary reasons are given.
6. Inability to reason about risks and benefits, even when secondary rationales are provided.
7. Inability to reach a reasonable conclusion, judged by standards of rationality.

These criteria focus on evaluating the *capacity to decide*, not whether a person's choice is socially acceptable or in line with mainstream values. A decision that appears irrational is not, by itself, grounds for deeming someone incapable.

Ultimately, evaluating a patient's capacity highlights that informed consent is a complex and multifaceted process. It depends not only on the individual's ability to understand information, reason through choices, and express preferences, but also on practical, social, and communication factors. Without careful attention to these challenges, traditional consent processes risk being ineffective or inconsistent. Recognizing these limitations provides the foundation for exploring more adaptive and personalized approaches to consent, which aim to support patients in making truly informed and voluntary decisions.

1.3 Background on AI

This section begins by examining the evolution of AI over the years, followed by an in-depth review of Generative IA and LLMs. We explore current Generative IA techniques, types, from basic NLP techniques to LLMs in Section 1.3.1. Then we explores more in detail how LLMs work and their different architectures in section 1.3.2. afterwards, we focus more on LLMs training methods in Section 1.3.2.4. Finally, we delve more in LLMs prompting and its' existing approaches in section 1.3.2.5.

1.3.1 Evolution of AI

Natural Language Processing (NLP) has undergone a remarkable transformation over the past few decades, evolving from simple statistical methods to sophisticated neural architectures capable of understanding and generating human-like text.

The evolution of statistical approaches in Natural Language Processing (NLP) began with fundamental techniques that laid the groundwork for modern language understanding systems. Early statistical methods demonstrated varying effectiveness across different text categorization algorithms. The Bag of Words model emerged as a foundational approach in the late 1990s, marking a crucial step in text categorization, handling high-dimensional feature spaces despite their computational simplicity. TF-IDF represented a significant advancement over basic BoW approaches. It improved retrieval effectiveness compared to simple binary term weights. After that, the transition to neural approaches in NLP marked a revolutionary shift in the field's capabilities. Deep learning methods have demonstrated remarkable improvements in NLP tasks, with neural networks achieving significant performance gains across multiple benchmarks while reducing the need for task-specific feature engineering. Indeed, early neural networks in NLP utilized feed-forward architectures that showed promising results but faced significant limitations. Deep learning approaches demonstrated the ability to learn word embeddings automatically, with models like word2vec representing words in dense vector spaces that captured semantic relationships. The emergence of RNNs, particularly Long Short-Term Memory (LSTM) networks, represented a significant advance in handling sequential data. LSTMs introduced gating mechanisms that allowed the network to selectively remember or forget information, significantly improving the handling of long-range dependencies. The introduction of the Transformer architecture in 2017 fundamentally changed NLP capabilities. The Transformer architecture's combination of parallelizable computation, effective handling of long-range dependencies, and state-of-the-art performance across benchmarks established it as the foundation for virtually all subsequent advances in NLP. Its ability to scale efficiently with computational resources while maintaining or improving performance paved the way for the development of increasingly powerful language models, marking a clear inflection point in the field's evolution, which led to rise of Large Language Models (LLMs) [26].

1.3.2 Generative AI and Language Models

In recent years, one of the most advances in Natural Language Processing (NLP) has been the development of Large Language Models (LLMs). These models, such as GPT, have shown capabilities in understanding and generating text. This progress has marked a new era for NLP, enabling systems that can communicate naturally and perform a wide variety of tasks with minimal supervision. Today, LLMs are viewed not just as tools for NLP, but as general-purpose foundation models with the potential to tackle diverse AI problems. In the following sections, we explore the core ideas behind LLMs and how they are built, scaled, and optimized.

1.3.2.1 Learning methods

Recent works and advancements in NLP has reached a high level with the emergence of foundation models, eg : transformers [27]. These advances are largely motivated by pre-training. Pre-training is the action of separating common components from neural network systems, and then train them on new data using self-supervision. This concept has been already used in deep learning with RNNs, autoencoders, and others [28]. But lately, it is used for various word embedding models deriving from BERT to GPT, following the same principle.

We note that a text embedding can be defined as a numerical dense representation of a word, phrase, sentence, or larger utterance in natural language in a certain space where texts with similar meanings are near each other. The meaning of a word is influenced by its context, and it is from this context that a word embedding is usually learnt. The meaning of a sentence is more complex because it depends on the words used in the sentence, the syntactic structure as well as the surrounding sentences [29].

The pre-training discussion in NLP have two types of problems : sequence modeling and sequence generation (encoding). Two different problems but can be defined with a single model as follows :

$$o = g(x_0, x_1, \ldots, x_m; \theta) = g_\theta(x_0, x_1, \ldots, x_m) \tag{1.1}$$

Here, $\{x_0, x_1, \ldots, x_m\}$ is a sequence of input tokens - Tokens can be thought of as pieces of words (units of text separated through tokenization) -, where x_0 is a special token (e.g., `[CLS]` or `<s>`) used to denote the start of the sequence. The function $g_\theta(\cdot)$ is a neural network parameterized by θ, and o denotes the output generated by this model.

The output o varies depending on the task:

- In *sequence modeling (token prediction)* tasks, o represents a probability distribution over the vocabulary.

- In *sequence encoding* tasks, o is a vector representation (embedding) of the entire input sequence.

The pre-training process addresses two major challenges:

- **Optimizing the model parameters θ for a general-purpose pre-training task.** Unlike traditional NLP models that are trained for specific tasks, pre-trained language models are designed to learn representations that are broadly applicable across multiple downstream tasks.
- **Transferring the pre-trained model $g_{\hat{\theta}}$ to downstream tasks.** This step involves adapting the pre-trained model to new tasks by either fine-tuning the parameters $\hat{\theta}$ using labeled data or by prompting the model with task-specific instructions.

In regards to pre-training, we differentiate between three basic approaches :

- Unsupervised learning : It was used in the early attempts of achieving pre-training, and it focuses on optimizing the parameters of a neural network using a criterion that is not directly related to specific tasks [30]. It is generally used as a preliminary step before supervised learning, helping in the dicovery of local minima and regularization in the training process.
- Supervised learning : In this approach, the model is combined with a classification layer (sentence classification for example). Then, the sequence model is adapted to downstream tasks. Essentially, the parameters are fine-tuned (adapted) using task-specific labeled data. The pre-training and fine-tuning of such model is straightforward, but requires more labeled data as the model grows [31].
- Self-supervised learning : In this case, the neural network generates it's own supervision signals. Instead of relying on initial model annotating the data, it trains from scratch and creates it's own signals from text. This is widely used in training sequence models.

1.3.2.2 Self-supervised pre-training approaches

To understand more about sequence encoding and generation, we need know about the different self-supervised architectures : decoder-only, encoder-only and encoder-decoder architectures. We focus basically on Transformer models, since they are the basis of most used pre-trained models in NLP [32].

a. Decoder-only architecture

Decoder-only architectures, like those used in GPT models [33], are widely adopted for language modeling. These models are typically implemented using Transformer decoders, with the cross-attention layers removed.

Given a sequence of tokens $x = \{x_0, x_1, ..., x_m\}$, a decoder $Decoder_\theta(\cdot)$ with parameters θ predicts the distribution of the next token x_{i+1} at each position i, conditioned on all previous tokens $\{x_0, ..., x_i\}$. The output is:

$$p_{i+1}^\theta = Pr_\theta(x_{i+1}|x_0, ..., x_i)$$

When training the model, and when it's complete, the model $Decoder_{\hat{\theta}}(\cdot)$ can be used to generate or evaluate text by computing probabilities of each token based on the preceding context[4].

b. Encoder-Only architecture

An encoder $Encoder_\theta(\cdot)$ is a function that takes a sequence of tokens $x = \{x_0, ..., x_m\}$ and produces a sequence of hidden vectors $H = \{h_0, ..., h_m\}$. Unlike decoder-based models, training encoder-only models is not straightforward because there is no obvious gold-standard output for the real-valued vectors.

A common solution is to couple the encoder with output layers that allow learning from easier-to-obtain supervision signals. In the pre-training phase, the encoder, together with a Softmax layer, is trained via self-supervision. In the application phase, the Softmax layer is removed, and the pre-trained encoder is combined with a prediction network to address specific problems.

$$\begin{bmatrix} p_1^{W,\theta} \\ \vdots \\ p_m^{W,\theta} \end{bmatrix} = \text{Softmax}_W(Encoder_\theta(x)) \tag{1.2}$$

Here, $p_i^{W,\theta}$ represents the predicted probability distribution at position i, i.e., $Pr(\cdot \mid x)$. The Softmax layer is parameterized by a matrix W, and we define it as:

$$\text{Softmax}_W(H) = \text{Softmax}(H \cdot W)$$

For brevity, we may omit the superscripts W and θ when referring to the probability distributions p_i.

[4]The term context refers to the preceding tokens in language modeling.

Difference from Decoder-Based Language Modeling. Although the model architecture in Equation 1.2 appears similar to that of decoder-based language models, there is a crucial difference in the semantics of the output.

In language modeling, p_i denotes the probability distribution of the next token x_{i+1} given the previous context $\{x_0, ..., x_i\}$. This follows an *auto-regressive* decoding pattern.

In contrast, encoder-based pre-training typically observes the entire sequence simultaneously. Consequently, there is no intrinsic directionality or "next word" prediction. Instead, tasks such as masked language modeling (e.g., BERT) or next sentence prediction are employed to generate supervision signals across the full input sequence [34].

c. Encoder-Decoder Pre-training

Encoder-decoder architectures are mainly employed to model sequence-to-sequence tasks such as machine translation and question answering. A simple yet powerful idea is to treat both the input and output as natural language texts, thereby allowing the encoder-decoder architecture to be applied in a unified *text-to-text* framework. For instance, given an input sentence, the model can be trained to generate a sentence describing its sentiment, such as *positive*, *negative*, or *neutral*.

This approach enables the development of a single text-to-text system capable of handling a variety of NLP tasks in a consistent manner. The process consists of two stages:

- **Pre-training:** A general-purpose encoder-decoder model is trained using self-supervised learning to acquire broad language understanding. This typically involves learning from large-scale unlabeled text corpora.
- **Fine-tuning:** The pre-trained model is subsequently adapted to specific downstream tasks by training it on task-specific, labeled text-to-text datasets.

By transforming problems into a common format, the encoder-decoder framework simplifies the architecture design and the multitask learning, like the T5 model [35].

1.3.2.3 Large Language Models

Language models are a manifestation of language modeling [36]. Their core objective is to estimate the probability of a sequence of tokens. Given a sequence $\{x_0, x_1, \ldots, x_m\}$, where x_0 is the start-of-sequence symbol $\langle \text{SOS} \rangle$, the joint probability of the entire sequence is computed using the chain rule of probability:

$$\Pr(x_0, \ldots, x_m) = \Pr(x_0) \cdot \Pr(x_1|x_0) \cdot \Pr(x_2|x_0, x_1) \cdots \Pr(x_m|x_0, \ldots, x_{m-1}) \tag{1.3}$$

This can be compactly written as:

$$\Pr(x_0, \ldots, x_m) = \prod_{i=0}^{m} \Pr(x_i | x_0, \ldots, x_{i-1}) \tag{1.4}$$

Or in logarithmic form:

$$\log \Pr(x_0, \ldots, x_m) = \sum_{i=0}^{m} \log \Pr(x_i | x_0, \ldots, x_{i-1}) \tag{1.5}$$

In both cases, $\Pr(x_i | x_0, \ldots, x_{i-1})$ denotes the probability of token x_i given all preceding tokens in the sequence.

In deep learning, we estimate these conditional probabilities using deep neural networks. A trained neural language model takes a sequence of tokens $x_0, \ldots, x_{i-1}$ as input and outputs a probability distribution over the vocabulary V, denoted as $\Pr(\cdot | x_0, \ldots, x_{i-1})$. The predicted probability for token x_i is simply the value of the i-th entry in this distribution.

An operation during inference is to select the most likely next token, given the context:

$$\hat{x}_i = \arg\max_{x_i \in V} \Pr(x_i | x_0, \ldots, x_{i-1}) \tag{1.6}$$

This token prediction can be applied iteratively to generate coherent text from left to right. Each newly predicted token $\hat{x}_i$ is appended to the context to predict the subsequent token $\hat{x}_{i+1}$, implementing the probabilistic formulations from Equations 1.4 and 1.5.

Example: Token Generation Process Table 1.1 illustrates how a language model generates a sequence of tokens given an initial prefix (context), here denoted by $\langle \text{SOS} \rangle$ a:

Context	**Prediction Rule**
$\langle \text{SOS} \rangle$ a	$\hat{x}_2 = \arg\max_{x_2 \in V} \Pr(x_2 \mid \langle \text{SOS} \rangle\ a)$
$\langle \text{SOS} \rangle$ a b	$\hat{x}_3 = \arg\max_{x_3 \in V} \Pr(x_3 \mid \langle \text{SOS} \rangle\ a\ b)$
$\langle \text{SOS} \rangle$ a b c	$\hat{x}_4 = \arg\max_{x_4 \in V} \Pr(x_4 \mid \langle \text{SOS} \rangle\ a\ b\ c)$

Table 1.1 – In each step, the model selects the token $x_i \in V$ that maximizes the conditional probability given the current context.

At each step, the model appends the most probable token to the current context and repeats the process. This illustrates the left-to-right generation characteristic of auto-regressive language models.

To summarize, LLMs operate by estimating the probability distribution over tokens con-

ditioned on prior context, and then generating text by selecting the most probable token.

Decoder-only Transformers Decoder-only Transformers are mostly used used in LLMs. These models are autoregressive: they generate text one token at a time, predicting each token based on all previous tokens in the sequence.

Input and Training The model takes a sequence of tokens $\{x_0, ..., x_{m-1}\}$ as input. At each position i, it computes a probability distribution $\Pr(\cdot \mid x_0, ..., x_{i-1})$ over the vocabulary, from which the next token x_i is sampled. The model is trained by maximizing the log-likelihood of the observed sequence:

$$\sum_{i=1}^{m} \log \Pr(x_i \mid x_0, ..., x_{i-1})$$

Embeddings Each token x_i is mapped to a d-dimensional vector e_i that combines both token and positional embeddings:

$$e_i = \text{TokenEmbed}(x_i) + \text{PositionalEmbed}(i)$$

Transformer Block The core of the model consists of a stack of Transformer blocks. Each block has two main components:

- A self-attention sub-layer
- A feed-forward network (FFN) sub-layer

Each sub-layer applies a transformation function $F(\cdot)$ using either a post-norm or pre-norm structure:

$$\text{Post-norm:} \quad \text{output} = \text{LNorm}(F(\text{input}) + \text{input})$$

$$\text{Pre-norm:} \quad \text{output} = \text{LNorm}(F(\text{input})) + \text{input}$$

Self-Attention Mechanism The self-attention mechanism operates using query, key, and value matrices Q, K, and V, derived from the hidden states:

$$\text{Att}_{\text{qkv}}(Q, K, V) = \text{Softmax}\left(\frac{QK^\top}{\sqrt{d}} + \text{Mask}\right)V$$

The `Mask` ensures that the model only attends to previous or current positions (no peeking ahead), where:

$$\text{Mask}_{i,k} = \begin{cases} 0 & \text{if } i \leq k \\ -\infty & \text{otherwise} \end{cases}$$

Multi-Head Attention Multiple attention heads process the input in parallel, each on a subspace of the hidden state:

$$Q^{[j]} = HW_q^j, \quad K^{[j]} = HW_k^j, \quad V^{[j]} = HW_v^j$$

$$\text{head}_j = \text{Att}_{\text{qkv}}(Q^{[j]}, K^{[j]}, V^{[j]})$$

$$F(H) = \text{Merge}(\text{head}_1, ..., \text{head}_\tau)W_{\text{head}}$$

Final Layer and Output After passing through L Transformer layers, a final linear projection and softmax layer compute the output token probabilities:

$$\begin{bmatrix} \Pr(\cdot \mid x_0) \\ \Pr(\cdot \mid x_0, x_1) \\ \vdots \\ \Pr(\cdot \mid x_0, ..., x_{m-1}) \end{bmatrix} = \text{Softmax}(H^L W_o)$$

where H^L is the output of the last Transformer block and W_o is the output projection matrix [27].

Generation Process The model generates text in an autoregressive manner:

$$\hat{x}_i = \arg\max_{x_i \in V} \Pr(x_i \mid x_0, ..., x_{i-1})$$

The predicted token $\hat{x}_i$ is appended to the context and used to predict the next token, continuing this process left-to-right.

Scalability These models are called *large* due to their deep architectures (many layers) and wide internal representations (large d and many attention heads), giving them powerful language modeling capabilities.

1.3.2.4 Training and Fine-tuning LLMs

a. Training LLMs To train an LLM, we start with a dataset D consisting of K token sequences. The model is trained using the "maximum likelihood estimation (MLE)" objective. For a sequence $x = x_0, \ldots, x_m$, the language model's log-likelihood is:

$$\mathcal{L}_\theta(x) = \sum_{i=1}^{m} \log \Pr_\theta(x_i \mid x_0, \ldots, x_{i-1})$$

The training goal is to find the model parameters $\hat{\theta}$ that maximize the total log-likelihood

over all sequences, which is optimized using standard gradient descent techniques :

$$\hat{\theta} = \arg\max_{\theta} \sum_{x \in D} \mathcal{L}_{\theta}(x)$$

Scaling and Challenges As models scale up in size and dataset volume, training becomes significantly more demanding. Yet larger models trained on more data tend to perform better [37], which has motivated the development of massive models.

However, this trend comes with several challenges:

- Infrastructure complexity: Training large models requires distributed computing systems to handle massive parameter counts and large datasets.
- Resource requirements: Training a model with tens of billions of parameters may require thousands of GPUs and considerable computing time.
- Training instability: Deep and wide networks may be prone to instability, requiring careful architectural and optimization adjustments.

b. Fine-tuning LLMs To overcome training challenges of LLMs, fine-tuning is introduced as a solution to perform specific NLP tasks, that generates text from task descriptions or prompts.

Let $x = x_0, \ldots, x_m$ be the context and $y = y_1, \ldots, y_n$ be the desired continuation. Fine-tuning involves learning to generate y given x:

$$\hat{y} = \arg\max_{y} \log \Pr(y \mid x) = \arg\max_{y} \sum_{i=1}^{n} \log \Pr(y_i \mid x_0, \ldots, x_m, y_1, \ldots, y_{i-1})$$

Instruction Fine-tuning and Zero-shot Generalization Fine-tuning helps activate the model's latent ability to follow instructions. Although LLMs learn general language patterns during pre-training, fine-tuning with instruction-labeled data helps the model to generalize across tasks, never seen before (with zero-shot and few-shot techniques) [38].

Fine-tuning Objective Let D_{tune} be the fine-tuning dataset, and let $\hat{\theta}$ be the parameters from pre-training. The fine-tuning goal is to adjust the parameters further:

$$\tilde{\theta} = \arg\max_{\hat{\theta}+} \sum_{\text{sample} \in D_{\text{tune}}} \mathcal{L}_{\hat{\theta}+}(\text{sample})$$

Each training sample consists of an input-output pair:

$$\text{sample} = [\text{ysample}, \text{xsample}]$$

The loss is computed only over the output portion:

$$\mathcal{L}_{\hat{\theta}+}(\text{sample}) = -\log \Pr_{\hat{\theta}+}(\text{ysample} \mid \text{xsample})$$

Backpropagation and Implementation During forward pass, the full sequence [ysample, xsample] is fed into the model. But during the backward pass, gradients are only propagated through the output tokens (ysample), ensuring only relevant parts of the network are updated.

Although cheaper than pre-training, fine-tuning still requires experimentation:

- Learning rate
- Batch size
- Number of fine-tuning steps

Several runs may be necessary to achieve satisfactory performance, making fine-tuning a non-trivial but critical phase in deploying LLMs effectively.

1.3.2.5 Prompting LLMs

Apart from using fine-tuning (due to the mentioned challenges), prompting emerges as an alternative solution. It involves supplying input text (the "prompt") that guides the LLM to generate a desired response. Prompts can range from simple instructions to rich structured contexts or conversations. Importantly, no additional training or fine-tuning is required.

Prompting allows us to move away from building specialized models for every task. Instead, we can use a single, well-trained model and adapt its behavior through prompts. As a result, prompt engineering has become a central topic in modern NLP.

One popular technique in prompting is to assign a 'role' to the LLM (e.g., acting as a psychologist). This role-playing enables LLMs to generate more focused, context-aware responses. An example can be a prompt that asks the LLM (in the role of a developmental psychologist) to explain 'delayed gratification'. The model generates a relevant response that reflects the assigned role.

However, not all tasks are easily handled through prompting. When complex reasoning is involved—such as arithmetic or commonsense reasoning—LLMs may still struggle. This has led to the development of **in-context learning (ICL)**, where demonstrations of how to solve a problem are included directly within the prompt. For instance, including a sample

math problem and its solution as part of the input can help the model learn how to solve a new, similar problem.

Yet, even ICL is not always sufficient. Solving multi-step reasoning tasks other strategies and techniques. In that regard, many prompt engineering techniques have emerged, spanning from zero-shot prompting to the latest advancements [39].

In the settings of our thesis, we focus more in the parts of reasoning and logic, and some basic techniques, as well as personalization, we explain that in the following paragraphs.

a. New Tasks Without Extensive Training

- **Zero-shot prompting** offers a paradigm shift in leveraging LLMs. This technique [40] removes the need for extensive training data, instead relying on carefully crafted prompts that guide the model toward novel tasks. Specifically, the model receives a task description in the prompt but lacks labeled data for training on specific input-output mappings.

```
{
  "prompt": "Explain the necessary materials, tools, and processes
  required to manufacture a car."
}
```

Listing 1.1 – Zero-Shot Prompting

- **Few-Shot Prompting** provides models with a few input-output examples to induce an understanding of a given task, unlike zero-shot prompting, where no examples are supplied [41]. Providing even a few high-quality examples has improved model performance on complex tasks compared to no demonstration. However, few-shot prompting requires additional tokens to include the examples, which may become prohibitive for longer text inputs. Moreover, the selection and composition of prompt examples can significantly influence model behavior, and biases like favoring frequent words may still affect few-shot results.

```
{
  "prompt": "Given these examples of manufacturing common-sense
  knowledge:
  1. A bicycle requires aluminum for its frame.
  2. Assembling a computer requires a clean, static-free
     environment.
  Describe the materials, tools, and assembly processes
  necessary for manufacturing a car."
}
```

Listing 1.2 – Few-Shot Prompting

b. Reasoning and Logic

- **Chain-of-Thought (CoT)** : Prompting LLMs often stumble in the face of complex reasoning, limiting their potential. Aiming to bridge this gap [11], introduced Chain-of-Thought (CoT) prompting as a technique to prompt LLMs in a way that facilitates coherent and step-by-step reasoning processes. The primary contribution lies in the proposal and exploration of CoT prompting, demonstrating its effectiveness in eliciting more structured and thoughtful responses from LLMs compared to traditional prompts. Through a series of experiments, the authors showcase the distinctive qualities of CoT prompting, emphasizing its ability to guide LLMs through a logical reasoning chain. This results in responses that reflect a deeper understanding of the given prompts. For example, the prompt would show the reasoning process and final answer for a multi-step math word problem and mimic how humans break down problems into logical intermediate steps. The authors achieved stateof-the-art performance in math and commonsense reasoning benchmarks by utilizing CoT prompts for PaLM 540B, achieving an accuracy of 90.2%.

```
{
  "prompt": "To understand what materials are necessary to build a
  car, first consider its main components. The body requires
  metal, the engine requires precision-machined parts, and
  the interior requires fabric or leather. Based on this,
  what tools are needed and how are these materials assembled?"
}
```

Listing 1.3 – Chain-of-Thought Prompting

- **Automatic Chain-of-Thought (Auto-CoT)** : [42] introduced Auto-CoT to automatically instruct LLMs with a "Let's think step-by-step" prompt to generate reasoning chains. Recognizing the possibility of errors in individually generated chains, AutoCoT enhances robustness through diverse sampling. It samples various questions and generates multiple distinct reasoning chains for each, forming a final set of demonstrations. This automated diverse sampling minimizes errors and enhances fewshot learning, eliminating the need for labor-intensive manual creation of reasoning chains. Auto-CoT demonstrated enhanced performance, surpassing the CoT paradigm with average accuracy improvements of 1.33% and 1.5% on arithmetic and symbolic reasoning tasks, respectively, employing GPT-3.

```
{
  "prompt": "Automatically generate a step-by-step reasoning
  process to identify the materials, tools, and manufacturing
  steps required to build a car, then provide the final answer."
}
```

Listing 1.4 – Automatic Chain-of-Thought Prompting

- **Self-Consistency** [43] introduced self-consistency, a decoding strategy enhancing reasoning performance compared to greedy decoding in CoT prompting. For complex reasoning tasks with multiple valid paths, self-consistency generates diverse reasoning chains by sampling from the language model's decoder. It then identifies the most consistent final answer by marginalizing these sampled chains. This approach capitalizes on the observation that problems requiring thoughtful analysis often entail greater

reasoning diversity, leading to a solution. The combination of self-consistency and chain-of-thought prompting results in significant accuracy improvements across various benchmarks, such as 17.9% on GSM8K, 11.0% on SVAMP, 12.2% on AQuA, 6.4% on StrategyQA, and 3.9% on ARC-challenge compared to the baseline chain-of-thought prompting.

```
{
  "prompt": "Explain the materials, tools, and processes required
  to manufacture a car. Provide three independent reasoning
  paths and select the most consistent conclusion."
}
```

Listing 1.5 – Self-Consistency Prompting

- **Logical Chain-of-Thought (LogiCoT)** : [44] proposes a Logical Chain-of-Thought (LogiCoT) prompting, a neurosymbolic framework that leverages principles from symbolic logic to enhance reasoning in a coherent and structured manner. Specifically, LogiCoT applies the concept of reductio ad absurdum to verify each step of reasoning generated by the model and provide targeted feedback to revise incorrect steps. LogiCoT can reduce logical errors and hallucinations through a think-verify-revise loop. Experimenting with Vicuna-33b and GPT-4, the findings underscore LogiCoT's notable enhancement of reasoning abilities, exhibiting improvements of 0.16% and 1.42% on the GSM8K dataset and 3.15% and 2.75% on the AQuA dataset compared to CoT, respectively.

```
{
  "prompt": "Using formal logical reasoning, determine the
  materials, tools, and processes required to manufacture
  a car. Justify each step using cause-effect relationships."
}
```

Listing 1.6 – Logical Chain-of-Thought Prompting

- **Chain-of-Symbol (CoS)** : [45] introduced CoS, employing condensed symbols instead of natural language. CoS provides distinct advantages: clear and concise prompts, heightened spatial reasoning for LLMs, and improved human interpretability. CoS suffers from challenges such as scalability, generalizab tegration with other techniques, and interpretability of LLM reasoning based on symbols. Notably, the implementation of CoS significantly elevates ChatGPT's performance, boosting accuracy from 31.8% to an impressive 92.6% on Brick World tasks. Moreover, CoS achieves up to a 65.8% reduction in prompt tokens, streamlining the process while maintaining high accuracy.

```
{
  "prompt": "Represent car manufacturing using symbolic
  abstractions such as:
  Material -> Component -> Assembly -> Product.
  Use this structure to infer the required materials,
  tools, and processes."
}
```

Listing 1.7 – Chain-of-Symbol Prompting

- **Tree-of-Thoughts (ToT)** : [46] proposed the Tree-ofThoughts (ToT) framework to enhance prompting capabilities for complex tasks requiring exploration and look-ahead reasoning. ToT extends CoT prompting by managing a tree structure of intermediate reasoning steps, known as "thoughts". Each thought represents a coherent language sequence moving toward the final solution. This structure allows language models to deliberately reason by assessing the progress generated by thoughts in solving the problem. ToT integrates the model's abilities to produce and evaluate thoughts with search algorithms like breadth-first or depth-first search. This enables systematic exploration among reasoning chains, with a look-ahead to expand promising directions and to backtrack when solutions are incorrect. ToT excelled in the Game of 24 tasks, achieving a 74% success rate compared to CoT's 4%. Additionally, in word-level tasks, ToT outperformed CoT with a 60% success rate versus 16

```
{
  "prompt": "Explore multiple reasoning branches for manufacturing
  a car by considering different material choices, tools,
  and assembly strategies, then select the best overall
  solution."
}
```

Listing 1.8 – Tree-of-Thoughts Prompting

- **Graph-of-Thoughts (GoT)** : [47] introduced the "Graph of Thoughts" prompting, a graph-based framework advancing traditional sequential methods to better align with the non-linear characteristics of human thinking. This framework permits dynamic interplay, backtracking, and evaluation of ideas, allowing the aggregation and combination of thoughts from various branches, departing from the linear structure of the tree of thoughts. The key contributions encompass modeling the reasoning process as a directed graph, offering a modular architecture with diverse transformation operations. The framework is presented as a versatile and dynamic approach to language model prompting, capturing the intricacies of human thought processes and enhancing model capabilities. The GoT reasoning model demonstrates substantial gains over the CoT baseline, improving accuracy by 3.41% with T5-base and 5.08% with T5-large on GSM8K. It also boosts accuracy over the state-of-the-art MultimodalCoT by 6.63% using T5-base and 1.09% with T5-large on ScienceQA.

- **System 2 Attention (S2A)** : [48] proposed System 2 Attention (S2A), utilizing the reasoning abilities of LLMs to selectively attend to relevant portions by regenerating the input context. S2A employs a two-step process to enhance attention and response quality by employing context regeneration and response generation with refined context. The effectiveness of S2A is evaluated across various tasks, including factual QA, long-form generation, and math word problems. In factual QA, S2A attains an accuracy of 80.3%, demonstrating a substantial enhancement in factuality. In long-form generation, it improves objectivity and receives a score of 3.82 out of 5.

```
{
  "prompt": "Deliberately analyze the problem of manufacturing a
  car using slow, step-by-step reasoning. Avoid intuitive
  shortcuts and explicitly verify each material, tool, and
```

```
  process before concluding."
}
```

Listing 1.9 – System 2 Attention Prompting

- **Thread of Thought (ThoT)** : [49] presented Thread of Thought (ThoT), a prompting technique designed to enhance the reasoning abilities of LLMs within chaotic contexts. ThoT, inspired by human cognition, systematically examines extensive contexts into manageable segments for incremental analysis, employing a two-phase approach where the LLM first summarizes and examines each segment before refining the information for a final response. ThoT's flexibility shines as a versatile "plug-andplay" module, enhancing reasoning across different models and prompting methods. Evaluations on question answering and conversation datasets reveal substantial performance improvements of 47.20% and 17.8%, respectively, especially in chaotic contexts.

- **Chain-of-Table** : The study by [50] introduced a pioneering prompting technique named Chain-of-Table. This method uses step-by-step tabular reasoning by dynamically generating and executing common SQL/- DataFrame operations on tables. The iterative nature of this process enhances intermediate results, empowering LLMs to make predictions through logically visualized reasoning chains. Significantly, Chain-of-Table consistently improves the performance of two benchmark tabular datasets by 8.69% on TabFact and 6.72% on WikiTQ, respectively.

```
{
  "prompt": "Construct a table incrementally describing materials,
  tools, and manufacturing processes for a car, and reason
  over the table entries to derive a complete description."
}
```

Listing 1.10 – Chain-of-Table Prompting

1.4 Knowledge Representation and Reasoning

Effective healthcare delivery and decision-making increasingly depend on the ability to represent, organize, and reason over complex medical knowledge. Modern healthcare systems generate vast amounts of heterogeneous data, from health records and pathologies to medical treatments. The challenge lies not only in managing this data but also in transforming it into actionable knowledge that supports clinical understanding and decision-making.

For this purpose, Knowledge representation provides formal frameworks to model this information, enabling interoperability, consistency, and advanced reasoning. In healthcare, it supports tasks such as clinical decision support, personalized treatment planning, predictive modeling, and patient-specific recommendations. By representing knowledge explicitly, through ontologies, knowledge graphs, and rule-based systems, healthcare applications can move beyond raw data processing to semantically rich reasoning that aligns with clinical knowledge and practice.

This section focuses on the main concepts of knowledge representation. We begin by defining ontologies and their construction in section 1.4.1. Then, we delve into the semantic web technologies and languages in section 1.4.2. Finally, we define knowledge graphs in section 1.4.3.

1.4.1 Ontologies

Ontologies play a central role today in the representation and modeling of knowledge. They allow domain knowledge to be formalized explicitly using formal languages, thereby facilitating its automated manipulation and sharing. The use of ontologies is widespread across various research fields, such as medicine [51], biology [52], among many others. In the medical domain in particular, ontologies provide a structured and interoperable way to represent clinical concepts, procedures, and relationships, supporting the integration of heterogeneous data sources and enhancing semantic reasoning. By establishing a shared conceptualization of domain knowledge, ontologies not only promote interoperability between systems but also enable advanced applications such as clinical decision support, personalized care, and automated knowledge discovery.

1.4.1.1 Ontology Definitions

The word "ontology" comes originally from philosophy. According to the *Oxford English Dictionary*, it refers to the study of being or existence in a general, abstract way. The *Merriam-Webster Dictionary* provides two meanings: one as a part of metaphysics that looks into the nature and relationships of being, and the other as a theory about the kinds of things that exist.

Nowadays, the term is used more loosely, especially in technical fields like computer science, information systems, and linguistics. It's often used to describe any structured way of organizing knowledge, especially when using Semantic Web technologies such as RDF (W3C).

In computer science, ontologies are widely used to describe a certain domain of interest using a formal knowledge representation. Many well-known ontology definitions have been proposed:

- "A specification of a representational vocabulary for a shared domain of discourse - definitions of classes, relations, functions, and others objects - is called an ontology." [53]
- "An ontology refers to an engineering artifact, constituted by a specific vocabulary used to describe a certain reality, plus a set of explicit assumptions regarding the intended meaning of the vocabulary words. This set of assumptions has usually the form of a first-order logical theory, where vocabulary words appear as unary or binary predicate names, respectively called concepts and relations. In the simplest case, an ontology describes a

hierarchy of concepts related by subsumption relationships; in more sophisticated cases, suitable axioms are added in order to express other relationships between concepts and to constrain their intended interpretation." [54]

- "An ontology is a formal, explicit specification of a shared conceptualization. A conceptualization refers to an abstract model of some phenomenon in the world by having identified the relevant concepts of that phenomenon. Explicit means that the type of concepts used, and the constraints on their use are explicitly defined. Formal refers to the fact that the ontology should be machine readable, which excludes natural language. Shared reflects the notion that an ontology captures consensual knowledge, that is, it is not private to some individual, but accepted by a group." [55]

- "An ontology specifies a shared understanding of a domain. It contains a set of generic concepts, together with their definitions and interrelationships." [56]

These definitions converge on several key aspects. An ontology represents shared knowledge of a domain, defined as a formal vocabulary of concepts, relations and axioms. IT enables the semantic annotation of data in a way that supports machine interpretation. It embodies the vision of a Web of linked data, promoted by the W3C. In the following sections, we outline the key concepts related to ontologies and their role in semantics.

According to [53], an ontology is formally defined as follows:

An **ontology** is a formal representation of knowledge that describes the key concepts of a domain, their relationships, and the rules or constraints that apply to them. An ontology O can be defined as a triple:

$$O = \langle C, R, A \rangle$$

where:

- C is the set of *classes* or concepts,

- R is a set of *relations* between those classes (typically expressing hierarchies or associations),

- A is a set of *axioms* that define logical rules about the classes and their relationships.

[57] points out that "ontology" has been applied to different types of classification systems and emphasizes the difference between ontologies, topic maps, and concept maps. While topic maps follow a standard and are structured, concept maps are more like diagrams used to show relationships visually. They also go beyond simple vocabularies or taxonomies by including logical reasoning capabilities. Ontologies are typically classified according to their level of abstraction [58] :

- **Glossary**: A flat list of terms with definitions, often lacking semantic relationships between terms.
- **Taxonomy**: A hierarchical classification that arranges concepts or entities from general to specific. It shows parent-child relations (e.g., Animal > Mammal > Dog).
- **Thesaurus**: A controlled vocabulary with defined relationships such as synonyms, broader or narrower terms, and related terms. [57] calls a thesaurus a "semantic road map" that improves search, indexing, and understanding across languages and fields.
- **Topic Map**: A semantic network that connects topics using standards for linking and representing knowledge. [57] notes that topic maps are structurally different from concept maps, which are typically visual and less formal.
- **Ontology**: The most formal and expressive of KOS types. In computer science, ontologies define entities, classes, and the relationships between them, often using standards like RDF. Ontologies support reasoning and can model constraints (e.g., a person can only marry another person).

This layered structure—from general to specific—helps ensure both reusability and precision in knowledge modeling. Ontologies are particularly valuable in semantic technologies and knowledge graphs, where formal representation and reasoning play a central role.

1.4.1.2 Ontology Construction

The construction of ontologies is a key step in knowledge engineering, as it enables the coherent structuring and organization of a domain's knowledge.

To assist domain experts in building ontologies, a methodology that guides the development pro- cess is required. In fact, an ontology development methodology comprises a set of established principles, processes, practices, methods, and activities used to design, construct, evaluate, and deploy ontologies [59]. Accordingly, different ontology development methodologies have emerged [60].

All of those methodologies divide the ontology development process into four main stages. The first step consists in identifying the objectives of the ontology. Next, the development step focuses on identifying and formally representing the concepts and relations within the domain. The third step, evaluation, assesses whether the resulting ontology satisfies the initial objectives. Finally, the documentation step ensures that the ontology can be easily reused and shared.

Several approaches can be used to build ontologies, each with its own advantages and limitations depending on the objectives, available resources, and application context. The main ontology construction approaches include building from scratch, reuse of existing ontologies, as well as methods leveraging large language models.

a) Building from scratch When ontologies are constructed *from scratch*, each author or team often proceeds in an ad hoc manner, without regard for broader alignment with other ontologies. As a consequence, such independently developed ontologies tend to suffer from poor interoperability with ontologies in related domains. In practice, even when mappings are later established between such ontologies, they are seldom maintained: as each ontology evolves on its own cycle, the mapping artifacts quickly become outdated, fragile, or inconsistent [61].

b) Reusing existing ontologies Reusing pre-existing ontologies or parts thereof offers significant advantages. Building on prior work reduces the time and effort required to model a domain and tends to increase ontology quality, since reused components have already been validated and exposed to usage, revealing and eliminating modeling errors. Furthermore, reuse supports easier maintenance over time, because shared components can be updated centrally or collaboratively, which reduces duplication of effort and fosters consistency across multiple ontologies [62, 63].

Specifically, reuse often relies on mid-level or upper-level ontologies to define classes and relations in a new ontology. This approach enhances semantic clarity, avoids ambiguity, and improves interoperability among ontologies grounded on the same foundational ontology. However, practical obstacles such as missing documentation, incompatible dependencies, or differing modeling choices can hinder reuse [64].

c) Using LLM-based approaches Recently, large language models (LLMs) have been proposed as tools to assist ontology engineering, for instance to suggest classes or relations, enrich existing ontologies, or help integrate textual definitions. Such approaches can reduce the manual burden on ontology engineers. Nonetheless, they are limited in highly specialized domains (e.g., biomedical or legal ontologies), where vocabulary, relations, and nuanced semantics may not be sufficiently represented in the training data. LLM-assisted ontologies therefore require careful validation to avoid introducing inconsistencies or ambiguous classes [65, 66].

1.4.2 Semantic Web technologies and languages

The Semantic Web, as defined by Tim Berners-Lee, transforms the traditional web of documents into a global web of data, where information is structured in a way that machines can understand and reason over [67]. Its goal is to make web content not only accessible to humans but also interpretable by software agents capable of performing complex tasks. To achieve this, technologies like the Resource Description Framework (RDF) and the Web Ontology Language (OWL) are employed. These tools enable the encoding of semantic meaning into data, facilitating interoperability and intelligent processing across the web. The Semantic Web, is defined as an extension of the current Web where information is given well-defined meaning, allowing computers and people to cooperate more effectively [68].

At the core of the Semantic Web lies a layered architecture "Semantic Web Stack" 1.1). This stack illustrates how various standards—from basic data formats to rules and logic—build upon each other to support the Web of Data. While the lower layers of the stack are already standardized (including RDF and OWL), the upper layers are still evolving. Ontologies, standardized by the W3C, play a critical role in this architecture by providing structured vocabularies that help link and contextualize information across the Web of Linked Data.

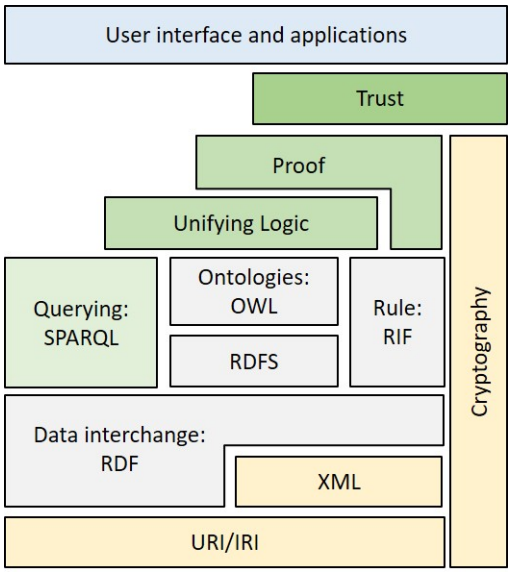


Figure 1.1 – Tim Berners-Lee's Semantic Web Stack [68].

To build applications that understand data on the Web, several technologies standardized by the W3C are used. This section briefly introduces the main ones relevant to this thesis.

1.4.2.1 Resource Description Framework (RDF)

RDF is a standard model for sharing and connecting data on the Web [69]. It represents information using triples: a subject, a predicate, and an object. For example, the sentence "The book has the title Semantic Web for the Working Ontologist" becomes a triple where "book" is the subject, "has title" is the predicate, and the title itself is the object (Figure 1.2).

Définition 1.1
RDF triple *Assume there are pairwise disjoint infinite sets I, B, and L (IRIs, Blank nodes, and literals). A triple* $(s, p, o) \in (I \cup B) \times I \times (I \cup B \cup L)$ *is called an RDF triple, where s is the subject, p the predicate, and o the object.*

The example in Figure 1.2 shows a representation of an RDF triple; the statement S "The book has the title Semantic Web for the Working Ontologist" is modeled as a triple with "book" as the subject, "has title" as the predicate, and "Semantic Web for the Working

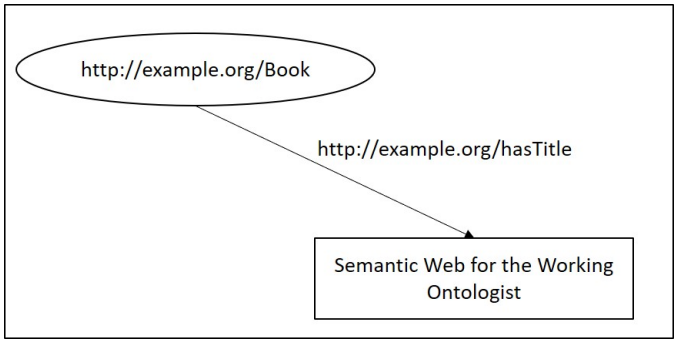


Figure 1.2 – Graph representation of an RDF triple.

Ontologist" as the object. Both the subject and the predicate are identified by a URI, while the object can be a URI or a literal value (i.e., a string or a number).

1.4.2.2 Web Ontology Language (OWL)

OWL extends RDF to describe richer relationships between concepts [70]. It allows defining classes, properties, and logical constraints. OWL comes in three versions:

- **OWL Lite**: for simple hierarchies,
- **OWL DL**: for more expressive but logically decidable models,
- **OWL Full**: the most expressive, but without computational guarantees.

OWL 2 further extends the language while staying compatible with RDF.

1.4.2.3 Differences and Usage of Ontology Formats

RDF, RDFS, and OWL are widely used ontology representation languages, but they differ in **expressivity, reasoning capabilities, and applications** [71, 72, 73]:

- **RDF (Resource Description Framework)** [74]:
  - Represents data as *triples* (subject, predicate, object).
  - Suitable for simple knowledge graphs, linked data, and basic semantic annotation.
  - Lightweight and web-friendly, but limited in expressing constraints or logic.
- **RDFS (RDF Schema)** [75]:
  - Extends RDF with class hierarchies and property definitions.
  - Useful for defining taxonomies and simple semantic relationships.

  - Supports basic hierarchical reasoning, but cannot encode complex constraints or cardinalities.

- **OWL (Web Ontology Language)** [76, 73]:
  - Builds on RDF/RDFS with high expressivity.
  - Enables rich reasoning: class equivalence, disjointness, property characteristics, and complex axioms.
  - Ideal for domains requiring automated reasoning, compliance checks, and advanced knowledge modeling (e.g., healthcare, law, biology).
  - Trade-off: more complex and computationally heavier.

Usage:

- Use **RDF** for lightweight linked data and simple annotations.
- Use **RDFS** when a structured taxonomy is needed without complex reasoning.
- Use **OWL** for domains requiring rich semantics, inference, and formal validation.
- Serialization formats (Turtle, JSON-LD, RDF/XML) are chosen based on context: web integration, readability, or computational efficiency [72].

1.4.2.4 SPARQL Protocol and RDF Query Language (SPARQL)

SPARQL is the query language used to retrieve and manipulate RDF data [77]. It supports various query forms like `SELECT`, `CONSTRUCT`, `ASK`, and `DESCRIBE`. A typical query includes namespace prefixes, dataset location, a pattern to match, and optional modifiers like sorting or limiting results.

Example:

```
PREFIX flight: <http://www.example.com#>
SELECT DISTINCT ?destination
FROM <http://www.w3.org/Flights/>
WHERE {
  ?f flight:DestinationCity ?destination.
  ?f flight:hasAirline ?airline.
  ?airline flight:hasName "AA";
}
ORDER BY(?destination)
```

Listing 1.11 – A simple SPARQL SELECT query.

1.4.2.5 SWRL rules

The Semantic Web Rule Language (SWRL) was designed to be the standard rule language of the Semantic Web [78]. It allows users to write rules expressed in terms of OWL concepts to reason with OWL individuals. The rules can be used to infer new knowledge from existing OWL knowledge bases.

To overcome the expressiveness limits of OWL2 (the newer version of OWL), we use the SWRL language. An SWRL rule includes two parts and is described as "antecedent" → "consequent." This signifies that if all the conditions in the antecedent are held, then all atoms in consequent must also be held. Atom is the basic component that appears in a SWRL rule. In SWRL, properties and individuals defined in the OWL are applied in the atom clause as the attribute and the parameter of the atom, respectively. There are many sorts of atoms, but in our work, two common atoms in SWRL syntax are used in the reasoning phase of problem-solving:

- C(?x): If x is an instance of the class C or the value of its data property, then C(?x) holds.
- P(?x, ?y): If x is related to y via property P, then P(?x, ?y) holds. Here P is the property defined in the existing ontology, x and y can be variables, individuals, or the data value.

It's to note that there are also built-in functions in the SWRL syntax that are capable of describing the logical comparison relationship, including *swrlb:lessThan* and *swrlb:greaterThan*.

1.4.2.6 Ontology programming tools

Many tools have been developped to help edition, maintenance and evaluation of ontologies, throughout the life-cycle of ontology creation [79]. the Protégé editor [80] is an open source tool that is widely used in this aspect.

As for ontology evaluation, some tools help ensuring the consistency and quality of the ontology, i.e., [81] and [82] reasoners. And since ontologies are mainly used as knowledge bases, some frameworks have been developed to assist semantic-based applications development, such as OWL API [83], OWLReady [84], Jena, RDF4J, etc.

1.4.3 Knowledge graphs

Although knowledge graphs (KGs) have existed and been utilized since the 1970s, their prominence has grown significantly since 2012, following the launch of Google's KG. Knowledge graphs are designed to represent information explicitly and are frequently employed in applications that require access to, or integration of, data from multiple sources. Their graph-based

structure provides a flexible and scalable way to model data, where nodes correspond to entities and edges capture the relationships between them. It usually relies on RDF and expresses facts as triples (subject, predicate, object).

The term *knowledge graph* is more than just a graph of connected data; it has a well-defined structure and purpose.

In a knowledge graph:

- Each **node** represents a concept or entity (like a person, place, or object). It has a unique identifier, typically in a machine-readable format, and may also include a human-readable label.
- Each **edge** represents a relationship between two nodes (e.g., "Alice knows Bob").

The key idea is that when multiple entities (nodes) share connections (edges), they form a rich semantic network. This network enables powerful reasoning, data integration, and knowledge discovery—especially when grounded in an ontology that defines the meaning of terms and their relationships.

knowledge graphs and ontologies have found applications across a wide range of domains due to their ability to structure and represent knowledge explicitly.

For example, iIn medicine, they support clinical decision-making, patient data integration, and biomedical research.

Beyond these, ontologies and KGs are widely used in finance, e-commerce, education, and cultural heritage, helping to integrate heterogeneous data sources, improve search and recommendation systems, and support intelligent reasoning over complex information. Their flexibility and formal representation make them valuable tools for both research and industrial applications.

1.5 Conclusion

In summary, this chapter has outlined the conceptual and theoretical foundations that underpin this research. It has shown that informed consent is not merely a legal or ethical formality but a dynamic process requiring adaptation to the complexity of modern healthcare and the diversity of patient needs. The analysis of its legal, ethical, and clinical dimensions revealed persistent challenges in ensuring clarity, personalization, and true patient engagement within current consent practices.

At the same time, the rapid development of Artificial Intelligence—particularly through Large Language Models and knowledge-based systems offers promising pathways to address these challenges. By combining the reasoning capacity of semantic technologies with the

generative capabilities of language models, it becomes possible to create consent processes that are both compliant with regulations and personalized to individual patients.

Finally, the exploration of knowledge representation approaches, including ontologies and knowledge graphs, highlights their central role in structuring medical information and enabling trustworthy AI integration. These technologies form the semantic layer necessary for interpretable, auditable, and context-aware personalization.

Building on these conceptual insights, Chapter 2 moves from theory to practice. It presents the overall system architecture developed in this thesis, detailing how the proposed framework operationalizes these principles to ensure legally compliant and personalized consent generation.

Chapter 2

SPIRIT : Semantic Personalized and complIant system for Informed consent and care RecommendaTion

Sommaire

2.1 Introduction 42

2.2 Semantic Personalized and complIant system for Informed consent and care RecommendaTion (SPIRIT) 42

2.2.1 Related works 42

2.2.2 SPIRIT architecture 45

2.2.2.1 Data acquisition and processing: 45

2.2.2.2 Personalized e-consent : 47

2.2.2.3 Medical recommendation : 48

2.2.3 Synthesis 49

2.3 Ontology-Driven machine learning annotation for Legal Compliance 49

2.3.1 Related works 50

2.3.1.1 NLP-based document analysis methods 50

2.3.1.2 Ontology-driven document representation 52

2.3.2 Ontology-Driven annotation architecture 54

2.3.2.1 Data Collection 55

2.3.2.2 Similarity and Annotation for Legal articles and consent Documents 56

2.3.3 Experimental results 59

2.3.3.1 Implementation Process 60

2.3.3.2 Experimental Setup 61

2.3.3.3 Parameter selection fo similarity algorithms 62

2.3.4 Knowledge graph embedding and annotation 62

2.3.5 Synthesis 64

2.4 Conclusion 65

2.1 Introduction

Following the conceptual foundations outlined in the precedent chapter, particularly the need to align legal, ethical, and clinical dimensions of informed consent, this chapter introduces an approach designed to operationalize these principles. In this work, we present SPIRIT, our proposed approach, in which, the combination of semantic reasoning and generative AI plays a central role in putting these principles into practice. SPIRIT integrates ontology-based reasoning with language model capabilities to ensure that consent documents can be automatically analyzed, personalized, and validated according to legal and ethical standards.

This chapter is organized into two main parts. The first part 2.2 situates the proposed approach within existing research and conceptual frameworks, offering an overview of current consent management solutions and positioning SPIRIT as an integrative, knowledge-driven architecture. The second part 2.3 introduces the first contribution of the thesis, an ontology-driven machine learning approach designed to ensure legal compliance through intelligent document annotation. This contribution focuses on the use of ontology-enhanced machine learning to achieve compliance-aware document understanding and annotation.

2.2 Semantic Personalized and complIant system for Informed consent and care RecommendaTion (SPIRIT)

This section introduces the proposed approach, SPIRIT, which serves as the foundation of this thesis. SPIRIT is designed to address the limitations of traditional consent processes by integrating semantic reasoning and generative AI to deliver consent experiences that are both personalized to patients and compliant with legal and ethical standards. Building on the multidimensional understanding of informed consent established in Chapter 1.2, this section conceptualizes consent as an evolving process involving medical, ethical, and legal knowledge.

The section first situates SPIRIT within existing research by reviewing state-of-the-art technological solutions and assessing their ability to meet the requirements of transparency, patient-centered, and compliance in Section 2.2.1. It then introduces the global architecture of SPIRIT, detailing its layered design and the interactions between its different components in Section 2.2.2. Together, these elements form the structural and functional backbone of the personalized consent and care recommendation approach proposed in this thesis.

2.2.1 Related works

As applied researchers, our state-of-the-art review has a dual focus. Firstly, it explores existing solutions in the French market for informed consent in the paramedical field. Simultaneously, it examines the issue of informed and voluntary consent in the healthcare system. Legislation and healthcare practices have accommodated these principles adeptly, as outlined in the

Kouchner law (2002) [1]. However, implementing them has proven challenging, particularly in light of the mentioned difficulties.

Today, with the advancement of technology and the need for digitalization, several companies offer solutions to meet the demand for consent without compromising patient autonomy. We explored the French market to identify the latest trends addressing this issue. *Je consens (Tessi)* [2]: is a solution focused on obtaining informed consent in the medical field. Designed for any healthcare institution, it offers an informative and interactive approach. Patients receive detailed information sheets and have the opportunity to participate in question-and-answer sessions with healthcare professionals to clarify all aspects of their treatment or intervention. This solution aims to ensure that patients fully understand the medical procedures to which they consent.

Easy consent (Calimed sante) [3], transforms medical documents into interactive questionnaires, accompanied by reformulations and visual illustrations. Through a video demonstration, *Easy consent* shows how to make the consent process more understandable for patients using visual aids and accessible wording. *Ordoclic* [4], offers digitalization of documents and enables secure document transmission by SMS and email, with the ability to collect electronic signatures for the returned document. This solution simplifies administrative management by eliminating constraints related to physical documents and automating signature processes.

Docaposte [5], solution relies on online signatures and digital archiving. By allowing electronic signature of consents and secure digital storage of medical documents. Docaposte facilitates future access to medical information and reduces dependence on paper documents.

Mind health (Frontline Media): define the features of the Docaposte solution. While specific details are limited, this solution appears to offer a similar approach tailored to the specific needs of healthcare institutions to enhance the efficiency of consent and medical document management.

Table 2.1 presents the different solutions and approaches regarding informed consent. Upon analyzing these solutions, we observed that each of them emphasizes digitization and consent collection but fails to specify their approach to effectively communicate with patients. Additionally, they seem to overlook the varying capabilities of patients based on their medical conditions, level of understanding, literacy, and more. Thus, the need for informed and voluntary consent remains unaddressed as a comprehensive solution in the market. However, in terms of digitization, they offer a promising foundation that we believe requires a semantic layer from both a modeling and substantive perspective.

The informed consent rely both on the patient's pathology, on his personal and social capabilities, environment, economical features, etc, And all the determinants described by

[1]Kouchner law : "LOI n° 2002-303 du 4 mars 2002 relative aux droits des malades et à la qualité du système de santé (1)"

[2]https://www.tessi.eu/fr/

[3]https://sante.easy-consent.fr/

[4]https://www.ordoclic.fr/

[5]https://www.docaposte.com/

Solution / Approach	**Main Features**	**Limitations / Gaps**	**Focus Area**
Je consens (Tessi)	• Interactive platform for medical consent • provides detailed information sheets and Q&A sessions with healthcare professionals.	• Focused on document management • lacks semantic understanding and adaptation to patient literacy or pathology.	Digital consent collection
Easy Consent (Calimed Santé)	• Transforms consent forms into interactive questionnaires • visual aids and reformulations for better understanding.	• Limited personalization • adaptation not context-aware (no patient profile integration or reasoning).	Patient understanding via multimedia
Ordoclic	• Digitizes and secures document transmission with e-signatures and centralized management.	• Emphasis on administrative efficiency rather than informed communication or personalization.	Document digitalization
Docaposte	• Enables electronic signatures and secure archiving of consent documents • suitable for multiple healthcare institutions.	• Does not handle patient comprehension or semantic linking to legal/medical context.	E-signature and secure storage
Mind Health (Frontline Media)	• Provides document management and consent tools for healthcare institutions.	• Similar to Docaposte • limited details and lack of semantic or personalized communication mechanisms.	Digital consent management

Table 2.1 – Related works on informed consent solutions

the WHO [85]. Moreover, the information system (SI) and the different types of data used; the complexity of the system makes it difficult to reach a consistent system able to handle these aspects.

Furthermore, working on healthcare system, reasoning provides more insights and understanding, which can be provided using knowledge graphs. This is the conclusion of the paper [86], showcasing the importance of knowledge graphs in intelligent medical application. However, we noticed that the proposed approaches suffer from a well-defined semantic connections between critical components, which hinders their capacity to communicate effectively with patients. Moreover, the available solutions do not holistically address the varying capabilities and needs of patients, overlooking the nuances introduced by distinct medical conditions, literacy levels, and other vital factors. The proposed approaches, despite their focus on digitization and consent collection, do not offer a comprehensive solution for addressing these issues. They lack of semantic layer, both in modeling and substantive aspects. Our aim is to propose a knowledge graphs and LLM-based approach to enhance semantic connections and provide more effective patient communication and adaptation within the medical e-consent process.

2.2.2 SPIRIT architecture

To address the challenges identified in the previous sections, this subsection presents the architecture of the proposed system SPIRIT. The core idea behind **SPIRIT** is to operationalize the notion of informed consent in its full legal, ethical, and medical complexity while enabling a personalized and comprehensible experience for each patient.

The system is designed to support both patients and healthcare professionals in the consent process by providing tools that ensure transparency, personalization, and compliance. Specifically, **SPIRIT** aims to establish a structured workflow that allows users to generate, adapt, and validate consent documents according to individual profiles, linguistic preferences, and medical conditions. Through this workflow, the system ensures that patients acquire a truly informed and contextualized understanding of the consent they provide. Figure 2.1 presents SPIRIT - the global architecture of our proposal -, which is composed of three main components: (1) Data acquisition and processing, (2) Personalized e-consent and (3) Medical recommendation.

2.2.2.1 Data acquisition and processing:

This component oversees core datasets and associated processing. Comprising two subcomponents, the first orchestrates offline mapping of consent documents to legal articles, ensuring validated forms and curating patient data for the knowledge graph. Detailed dataset information is provided in the following subsection. The second sub-component details access modalities and required processing for each dataset, focusing on format and data specifics.

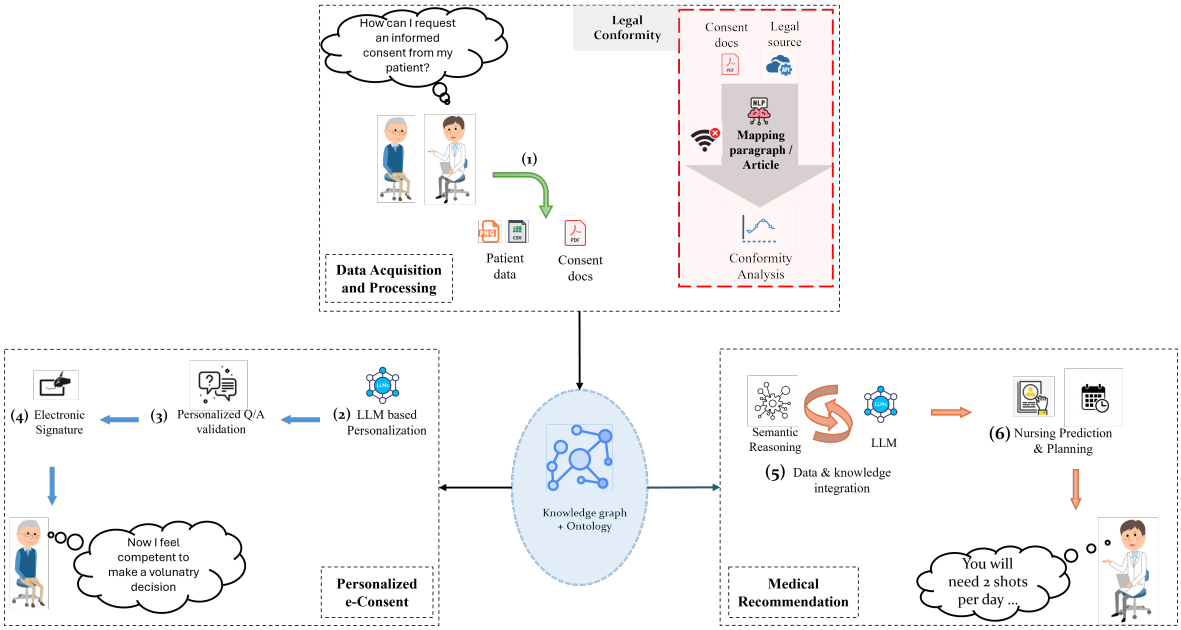


Figure 2.1 – SPIRIT : Semantic Personalized and Compliant System for Informed Consent and care Recommendation

a. Datasets description:

We present in this subsection the various datasets we work with and their acess process, with a focus on precision and granularity:

- **Patient Dataset:** This dataset encompasses personal details (names, addresses, contact info) and delves into each patient's medical profile, covering medical history, current conditions, medications and allergies. Data is maintained in structured CSV files. The *Dossier Médical Partagé* offers an alternative for compliant data access.

  - Access to patient data is achieved through a multi-modal approach: (1) CSV Files: patient data is primarily collected and maintained in structured CSV files, ensuring integration into the system. This mode of access enables efficient data retrieval and management; (2) Shared Medical Folder: the system securely interfaces with the Dossier Médical Partagé using APIs, ensuring up-to-date and compliant patient data. Access to the patient's dataset involves both structured and unstructured data. This structured data includes personal details, medical histories, allergies, and more. In contrast, unstructured data in the form of scanned PDF files contains handwritten notes, medical terminology, and diverse medical documents, all of which are processed within the system.

- **Consent Forms:** The consent forms presented to patients are detailed and categorized into various sections, containing textual descriptions. These forms come in two types: (1) Documents to Fill: patients are required to provide precise information in these

documents, such as personal information; (2) Documents to Read and Consent to: these documents offer detailed information about medical interventions with legal references.

- Access to consent documents is facilitated through Optical Character Recognition (OCR) technology, paper consent documents are scanned and digitized into PDF files, which are then directly accessed and processed within the system. The consent documents further exemplify the integration of structured and unstructured data. Traditional paper consent documents are scanned and digitized into structured PDF files, which are then directly accessed and processed within the system.

- **Légifrance API**[6] Employed for legal compliance, this dataset validates consent form content against legal requirements in real-time through secure API integration, ensuring legal compliance for each consent document.

  - The database is accessible through secure digital channels, the system initiates API calls to Légifrance, obtaining relevant legal articles in structured JSON format. The response is processed to ensure legal compliance.

The Access process ensures that all data access modalities are secure, compliant with data protection regulations, and seamlessly integrated into the e-consent approach. Patient data access adheres to strict privacy and consent protocols, safeguarding patient privacy and data security.

b. Mapping Paragraph/Article:

This step ensures legal validation of the consent documents based on the legal source. It is an offline process based on a similarity measure between the documents and legal articles. The access via API to the legal source allows the validation to be updated following each change of law. The result is textual legal consent forms, that are fed to the personalized e-consent component (this process will be detailed in section 2.3.2).

2.2.2.2 Personalized e-consent :

This component aims to personalize consent forms by leveraging LLMs to refine, rephrase, and summarize documents. This personalization is based on the patient's unique information stored in the knowledge graph, including literacy, language, medical history, and health condition. The result is customized document, upon which we generate questions and answers that assess the patient's comprehension. afterwards, we establish an electronic signature upon satisfactory responses, confirming well-informed consent.

This component contains three main steps : (a) LLM-based personalization, (b) Personalized Q&R and (c) Electronic signature.

[6]https://www.legifrance.gouv.fr/

a. LLM-based personalization

The purpose of step (2) is to personalize documents in accordance with the patient's profile. In this step, an LLM is used to personalize consent forms extracted from prior data acquisition. The process is based on LLM reformulation. The patient's data is extracted from the knowledge graph and included in the context given to the LLM. This gives parameters for the LLM generation (literacy level, language, pathology). In addition, the consent information is extracted from the documents and given as well to the context, which increases accuracy of data generation of the LLM.

b. Personalised Q&R

The purpose of step (3) is to assess the patient's comprehension regarding the text presented. In this step, the personalized Q&R module generates questions and answers, verifying the patient's comprehension of the information. Based on the previous content generated from the LLM, the system sets another context with the same parameters for each patient. This process is based on an LLM, aiming to generate the questions based on the content provided. Ultimately, the LLM generates choice questions to easily confirm the patient's comprehension.

c. Electronic signature

After evaluating the patient's understanding, Step (4) enables the establishment of an electronic signature. This signifies the patient's explicit consent and acknowledgment of presented information, reinforcing the legally binding nature of their agreement. It serves as a secure and traceable mechanism for the consent.

2.2.2.3 Medical recommendation :

The purpose of this component is to assist healthcare professionals in proposing personalized care plans aligned with patient profiles, clinical guidelines, and institutional service capabilities. This component transforms the knowledge and personalization layers converge into actionable recommendations for nursing care and service orientation. The component contains two main steps : (a) Data & knowledge integration and (b) Nursing prediction and planning.

a. Data & Knowledge Integration

Step (5) aims to integrate ontology guidelines with local patient's data using LLM based mappings. This semantic representation enables reasoning across both medical and contextual dimensions (e.g., living conditions, autonomy, chronic diseases). This step not only serves as

a matching between the local dataset and ontologies, but also has a parameter augmentation purpose, that is fed to the prediction model.

b. Nursing prediction and planning

The purpose of step (6) is to propose and recommend several nursing procedures for a given patient. This prediction is based on a multi-label classification, that takes as input the augmented data from the previous step, and generates then a list of recommended procedures, that will assist the professional in the care plan proposal.

2.2.3 Synthesis

In summary, this section presented **SPIRIT** — the Semantic Personalized and complIant system for Informed consent and care RecommendaTion. The proposed approach provides a unified framework that bridges the legal, ethical, and medical dimensions of informed consent through the integration of symbolic reasoning and generative AI.

Beyond legal compliance, SPIRIT explicitly addresses ethical requirements, including respect for patient autonomy, transparency of information, and fairness in access to care. By enabling personalized, understandable, and context-aware consent documents, the system aims to reduce informational asymmetries and support truly informed decision-making, while preserving the central role of healthcare professionals in the process.

By structuring the system into three main components—data acquisition and processing, personalized e-consent generation, and medical recommendation—SPIRIT establishes a foundation for developing transparent, patient-centered, ethically grounded, and legally compliant consent systems. This global approach lays the groundwork for the subsequent sections, where specific components of SPIRIT are detailed and evaluated.

2.3 Ontology-Driven machine learning annotation for Legal Compliance

This section presents the first contribution of this research: an ontology-driven machine learning approach dedicated to ensuring the legal compliance of consent documents. While **SPIRIT** provides the overall framework for integrating personalization and reasoning, this contribution focuses on the first pillar, specifically on verifying that the documents generated or adapted for patients respect the applicable legal and ethical constraints.

To establish this contribution, the section begins with a review on one hand, of NLP and embedding techniques commonly used for text understanding and document analysis, with a focus on their application to the healthcare and legal domains (Section 2.3.1.1). On the other

hand, a review of ontologies regarding document representation in the domain of healthcare (Section 2.3.1.2). It then introduces the proposed approach, **Ontology-Driven Machine Learning Annotation for Legal Compliance**, which leverages ontological knowledge to guide and enrich machine learning models in classifying and annotating consent documents (Section 2.3.2). Finally, the section presents the experimental setup and evaluation (Section 2.3.3), examining how different model configurations and semantic layers contribute to improving annotation accuracy and compliance verification.

2.3.1 Related works

Ensuring the legal compliance of consent documents requires methods capable of both understanding natural language content and interpreting its meaning within a structured legal and clinical context. Existing research in this area can be broadly divided into two complementary perspectives: NLP-based document analysis and ontology-driven document representation. The first focuses on extracting and classifying information from unstructured text using machine learning, embeddings, and transformer-based models, while the second emphasizes the formal modeling of document knowledge through ontologies to enable semantic reasoning and interoperability.

In the following subsections, we first review the main NLP methods and recent advances in document analysis (section 2.3.1.1), highlighting their contributions and limitations in capturing domain-specific semantics. We then examine existing ontologies for document representation (section 2.3.1.2), emphasizing their structure, use cases, and relevance to legal and clinical contexts. Together, these perspectives provide the foundation for our proposed ontology-driven annotation approach.

2.3.1.1 NLP-based document analysis methods

Document analysis plays a vital role in domains such as legal, healthcare, and governance, facilitating tasks like compliance checking, contract management, and decision-making. Natural Language Processing (NLP) has significantly enhanced document analysis by extracting insights from unstructured data, thereby improving scalability and efficiency. Traditional methods like Bag-of-Words (BoW) [87] and Term Frequency-Inverse Document Frequency (TF-IDF) [88] were foundational due to their simplicity. BoW represents documents as collections of word frequencies, while TF-IDF highlights informative terms based on their rarity across the corpus. However, these non-contextual methods fail to capture semantic relationships between words, limiting their effectiveness for nuanced tasks like document similarity assessment and annotation. This limitation has driven the transition toward context-aware approaches [89, 90, 91].

The evolution of NLP has introduced advanced techniques, such as word embeddings and sequence models, which better capture semantic and contextual relationships in language. Word2Vec, GloVe, and FastText represent words as dense vectors, encoding semantic sim-

ilarities based on context. Sequential models like recurrent neural networks (RNNs) and long short-term memory networks (LSTMs) further advanced NLP by retaining context over longer sequences of text. These methods outperform traditional approaches by capturing the meaning of words within their broader linguistic context, offering deeper and more nuanced document analysis capabilities [92, 93].

The advent of transformer-based models has revolutionized the field, delivering state-of-the-art performance in document similarity and annotation tasks. Models such as BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pre-trained Transformer) represent significant advancements in understanding and processing text. BERT's bidirectional architecture captures the context of words from both preceding and succeeding text, excelling in tasks like similarity analysis, sentiment detection, and entity recognition. Conversely, GPT employs an autoregressive framework to generate coherent text, making it well-suited for language modeling and text generation tasks. These models leverage pre-trained architectures to encode documents into high-dimensional vector spaces, enabling precise similarity measurement through metrics like cosine similarity. Fine-tuning these models on domain-specific datasets has proven to be effective for specialized tasks such as cross-referencing legal and healthcare documents, as well as information extraction [94, 95, 96].

Transformer-based models' ability to encode contextual relationships and model long-range dependencies has made them indispensable for modern document analysis. Their application in tasks such as document similarity measurement, annotation, and classification continues to drive innovation, particularly in complex domains requiring nuanced textual understanding [96, 94].

In terms of document analysis, many works inspired from the mentioned methods have been deployed. For instance, Headword Attention-based Recurrent Convolutional Neural Network (HA-RCNN) combines RNN, CNN, and attention mechanisms to effectively capture contextual information and keywords in legal documents, achieving efficient results in document classification and similarity assessments [97]. Similarly, Supervised Multiple Positives and Negatives (SupMPN) leverages a supervised learning approach with multiple positive and negative examples to generate robust representations, leading to improved accuracy in document categorization and similarity measurement [98]. Hybrid methods such as the Hybrid Method for Sentence Similarity Calculation (HybMethod) combine word embeddings, lexical databases, and word order to compute similarity scores. This approach has been extensively used to compare clauses, case summaries, and other texts, enhancing document comparison and retrieval [99]. Short-Text Semantic Similarity (STSS) methods focus on evaluating semantic similarity in short texts like legal opinions and case summaries. By employing word embeddings and neural networks, these methods enable the identification of similar precedents, supporting decision-making in document analysis [97, 100].

Beyond BERT and GPT-based models, other transformer and embedding architectures have also influenced document similarity research. ELMo introduced deep contextualized word representations, modeling word meaning dynamically based on sentence-level context. Doc2Vec extended vector-based representations to entire documents, allowing direct document-level embeddings. RoBERTa and XLNet further advanced transformer architectures by opti-

mizing training procedures and employing permutation-based language modeling, achieving improved results in semantic similarity and sentence-pair classification. These models have become foundational baselines for legal, clinical, and contractual document similarity due to their ability to encode multi-sentence reasoning and contextual relations.

These advancements underscore the transformative role of NLP in document analysis, moving from traditional approaches to context-aware models such as HA-RCNN, SupMPN, HybMethod, and STSS. By leveraging semantic understanding and computational techniques, these methods have significantly enhanced operational efficiency and decision-making. Nonetheless, challenges such as handling domain-specific nuances and ensuring scalability remain open for future research. Table 2.2 represents a comparative analysis of the reviewed text similarity measures.

Synthesis. The comparative analysis highlights a clear evolution from statistical to semantic and contextual methods for document similarity measurement. Traditional approaches such as Word2Vec and FastText remain efficient for generating compact word embeddings but fail to capture complex syntactic and contextual nuances at the document level. Conversely, transformer-based models such as BERT, RoBERTa, and XLNet have demonstrated superior performance in understanding semantic relations and contextual dependencies across sentences, which is crucial in highly sensitive domains like healthcare and legal compliance. Given these strengths, our approach leverages both models with a layer of semantic presentation to ensure precise semantic alignment and compliance verification within informed consent documents. This methodological choice balances performance, adaptability to domain-specific data, and robustness.

2.3.1.2 Ontology-driven document representation

The representation and semantic analysis of documents in healthcare and legal contexts have been supported by various ontologies, each designed with specific goals and levels of abstraction. Ontologies provide the formal structure needed to describe document types, their components, and the relationships between content elements, enabling automated reasoning, information retrieval, and compliance verification.

In the biomedical domain, several ontologies and terminologies have been adopted to describe clinical content. SNOMED CT (Systematized Nomenclature of Medicine—Clinical Terms) [7] offers comprehensive coverage of clinical concepts, observations, and procedures, but its primary focus is on clinical events rather than document structures. ICD (International Classification of Diseases) [101] and ICF (International Classification of Functioning, Disability and Health) [102] serve as classification systems for diagnoses and functional assessments, supporting reporting and epidemiological analysis rather than document annotation. The UMLS Metathesaurus [103] integrates these vocabularies to facilitate interoperability and concept mapping, but it does not explicitly address document metadata or typology.

Beyond these clinical terminologies, several ontologies have been proposed to model the structural and semantic aspects of documents. The Document Ontology (DOCO) [104] pro-

Category	**Method**	**Strengths**	**Limitations**	**Typical Use in Legal/-Clinical Documents**
Statistical Methods	BoW / TF-IDF [87, 88]	Simple, interpretable, fast	No context or semantics	Keyword extraction, clause matching
	HybMethod [99]	Combines lexical + embedding	Requires manual features	Contract and case text comparison
Neural Embeddings	Word2Vec [92]	Good semantic similarity	Needs large corpus	Term/phrase similarity detection
	FastText [93]	Handles morphology/-subwords	Higher computation cost	Handling medical/legal terminology
	Doc2Vec [92]	Document-level vectors	Requires large training data	Document-level similarity
Contextual Embeddings	ELMo [94]	Dynamic contextual meaning	Resource intensive	Word-sense disambiguation
	BERT [96]	Bidirectional semantic modeling	Requires fine-tuning	Entity detection, semantic similarity
	RoBERTa [94]	Optimized training, robust	Requires GPU training	Legal/clinical semantic reasoning
	XLNet [94]	Long-range context modeling	Training complexity	Multi-sentence comparison
	GPT [95]	Coherent generation	Autoregressive limitations	Summarization, classification
Hybrid/Domain-Specific Methods	HA-RCNN [97]	Attention + CNN + RNN	Domain-specific architecture	Legal case classification
	SupMPN [98]	Uses multiple pos/neg samples	Needs supervised dataset	Clause similarity, contract matching
	STSS Methods [100]	Very strong for short texts	Limited long-text reasonining	Case summaries, precedent similarity

Table 2.2 – Comparative Analysis of Text Similarity Methods

vides a general framework for representing the logical structure of documents, including sections, paragraphs, and metadata, mainly in scholarly or administrative contexts. The FRBRoo (Functional Requirements for Bibliographic Records - object-oriented) [105] and the CIDOC [106] ontologies have also been used to represent complex documentation workflows, though their orientation remains cultural and archival rather than clinical.

In healthcare-specific document representation, the LOINC (Logical Observation Identifiers Names and Codes) [107] ontology stands out. LOINC was designed to standardize the identification of laboratory tests, clinical observations, and document types. It provides a robust typology for medical documents through axes such as *Kind of document*, *Type of service*, *Setting*, *Role*, and *Subject matter domain*. These axes make LOINC particularly well suited for describing clinical and administrative documents, supporting both semantic interoperability and reasoning over document metadata. Moreover, LOINC's integration with HL7 and FHIR standards reinforces its role as a bridge between structured data and document-level semantics in healthcare systems.

Synthesis. From this review, it appears that while several ontologies and terminologies contribute to the semantic representation of medical information, few directly address the structure and classification of clinical documents. General ontologies such as DOCO or FRBRoo provide valuable models for document organization but lack medical specificity. Conversely, clinical terminologies such as SNOMED CT, ICD, or UMLS excel at representing medical concepts yet do not capture the metadata necessary to describe document kinds, services, or contexts.

The LOINC ontology offers a balanced and interoperable solution by focusing on the semantics of medical and administrative documents. Furthermore, its alignment with major interoperability standards (HL7, FHIR) ensures that document annotations can be integrated into broader healthcare information systems. On top of that, Loinc includes concepts related to both regulations and informed consent, which are vital for the scope of this research.

For these reasons, LOINC was selected as the reference ontology in our approach. Its ability to represent the document layer—bridging legal, administrative, and clinical semantics—makes it particularly suited for the ontology-driven machine learning framework proposed in this thesis, which aims to ensure legal compliance and semantic coherence in consent document analysis.

2.3.2 Ontology-Driven annotation architecture

Building upon the legal compliance dimension of the SPIRIT framework, this subsection introduces the proposed *Ontology-Driven Annotation Architecture*. The goal of this approach is to ensure that consent documents comply with existing legal standards while preserving semantic interpretability for automated analysis.

As discussed in the previous sections, informed consent documents are complex legal entities that encode obligations, permissions, and prohibitions. To automatically verify their

conformity, it is necessary to move beyond purely syntactic document representations toward semantically enriched structures that reflect both legal and clinical meaning.

Our approach leverages the **LOINC ontology** as a semantic backbone for document modeling, given its structured taxonomy of document types, roles, and contexts. By aligning consent document components with LOINC classes and relations, we enable an annotation process that can later support reasoning and compliance verification.

The proposed architecture operates through a sequence of four main components: input, preprocessing, similarity computation, and output. First, raw consent documents are ingested as the input data. A preprocessing module performs text normalization operations such as filtering, tokenization, and lemmatization to ensure the quality and consistency of the textual corpus. The second component computes semantic representations of the processed text using machine-learning-based embedding models (e.g., Word2Vec, Doc2Vec, BERT, or XLNet). At this step, embeddings can be enriched through ontology-based alignment with domain knowledge from the consent knowledge graph, enabling a deeper semantic representation of clinical and legal terminology. The similarity measurement module then calculates document and paragraph similarity using the refined embeddings, allowing the system to detect relevant semantic matches between text fragments. Finally, the output component annotates similar segments of consent documents based on a predefined similarity threshold T and produces structured annotations. Overall, the architecture supports (1) similarity computation using advanced embeddings, (2) embedding enhancement via ontology-driven knowledge alignment, and (3) automated annotation of consent documents.

Together, these aspects form a pipeline (figure 2.2) capable of bridging natural language content and formal legal structures, thus facilitating the automated annotation of consent documents. The following subsections detail the workflow and its evaluation, including dataset preparation, hyperparameter optimization, and the impact of the ontology as a semantic layer.

2.3.2.1 Data Collection

The dataset used in this study comprises two main categories of textual sources: **legal articles** and **consent documents**. Legal articles refer to legislative and regulatory texts that define the legal framework governing informed consent in healthcare, such as national health laws, data protection regulations, and ethical guidelines. These documents establish the obligations, rights, and procedures that healthcare institutions must comply with when obtaining patient consent.

Conversely, consent documents represent the *operational* implementation of these legal principles within healthcare settings. They are the actual forms or digital interfaces presented to patients to obtain consent for medical procedures, data use, or care delivery. These documents reflect how legal requirements are translated into clinical and administrative practice.

Legal articles were retrieved via a dedicated API based on predefined criteria, including

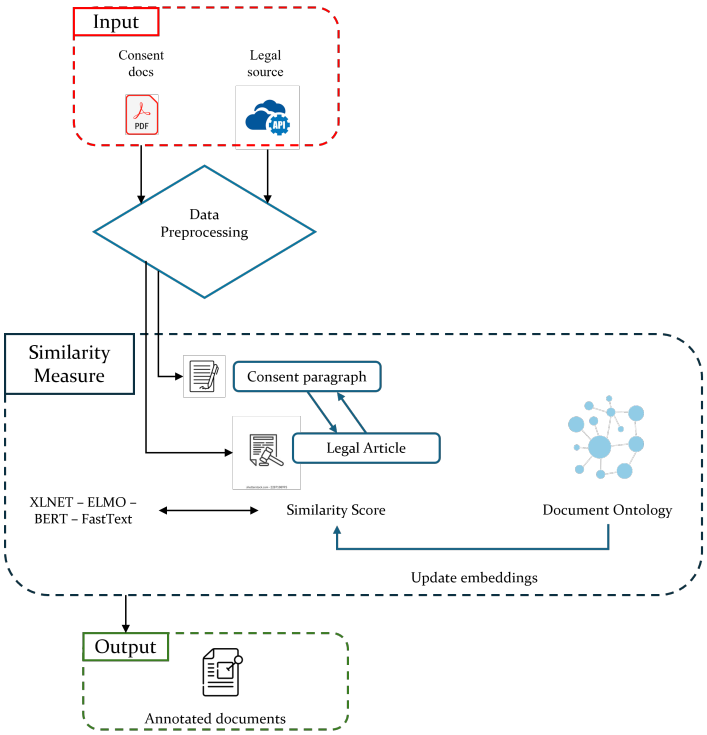


Figure 2.2 – Ontology-Driven machine learning annotation for Legal Compliance

domain-specific keywords, publication dates, and relevance to informed consent. For each document, both the full text and associated metadata—such as article titles, unique identifiers, sections, publication dates, legal status, and thematic categories—were extracted and stored for further processing. An illustrative example of the collected metadata is shown in Table 2.3. This process resulted in a curated corpus of legal and consent-related texts serving as the foundation for ontology-driven annotation and compliance analysis.

Simultaneously, consent documents provided by the Léonie Chaptal foundation were collected. These documents underwent some preprocessing steps, including noise removal, text extraction, and format standardization, ensuring data consistency and relevance for subsequent analysis.

To ensure semantic consistency across these heterogeneous sources, we employ the LOINC document ontology as the underlying reference model, as previously established in Section 2.3.1.2. Figure 2.3 shows an overview of Loinc DO axes, and the concepts related to consent and legal compliance we're interested in.

2.3.2.2 Similarity and Annotation for Legal articles and consent Documents

This section introduces algorithms for semantic similarity analysis and annotation, enhanced by integrating a knowledge graph (KG) to refine similarity scores using domain ontology. The process of how is detailed below :

Metadata Field	Example
Article Title	Code des postes et des communications électroniques
Article ID	LEGITEXT000006070987
Section Title	Section 1 : Autorité de régulation des communications électroniques, des postes et de la distribution de la presse
Article Number	L36-11
Legal Status	MODIFIE
Publication Date	2018-11-25

Table 2.3 – Legifrance API Response Metadata

- 1. Input:
  - Raw consent documents and domain knowledge from the knowledge graph.
- 2. Preprocessing:
  - Clean the dataset, tokenize, remove noise, and lemmatize the text.
- 3. Similarity Measurement:
  - Vectorize documents using embeddings (e.g., Word2Vec, XLNet).
  - Augment embeddings with knowledge-graph concepts and ontology alignment.
  - Compute semantic similarity and identify matches above a threshold T.
- 4. Output:
  - Automatically annotate the document segments that satisfy the similarity criteria.

a - Similarity Analysis

The similarity analysis calculates semantic alignment between legal articles and consent documents using text embedding models (e.g., Word2Vec, XLNet). Preprocessed documents are vectorized into embeddings, and cosine similarity is computed to measure alignment. Additionally, knowledge graph augmentation refines the similarity process by enriching embeddings with contextual relationships.

The preprocessing and vectorization steps are essential for preparing the dataset and extracting meaningful embeddings. During preprocessing, each document is cleaned by removing special characters and then tokenized and lemmatized to ensure a uniform and simplified representation. For vectorization, documents are embedded using models like Word2Vec, FastText, and XLNet. Key parameters, such as window sizes (ranging from 2 to 15) and vector dimensions (50, 100, 200, or 300), are adjusted to optimize the embeddings for capturing semantic and syntactic relationships. The result is a matrix of document embeddings that serves as the foundation for subsequent similarity computations and analyses.

Axis	Description
Kind of Document (KoD)	Specifies the structural format or generic type of the document.
Type of Service (ToS)	Describes the kind of healthcare service or encounter being documented.
Subject Matter Domain (SMD)	Specifies the clinical specialty or body system relevant to the document's content.
Role	Identifies the occupation or specific role of the person primarily responsible for the service or document's authorship.
Setting	Defines the physical location or channel where the clinical care was provided.

(a) LOINC Document Ontology Axes



(b) Concepts of interest visualised in Protégé

Figure 2.3 – Overview of the LOINC Document Ontology axes it's concepts.

b - Knowledge graph-Enhanced Similarity Analysis

The Knowledge graph-Enhanced Similarity algorithm leverages annotated documents and domain knowledge stored in a knowledge graph. It refines embeddings with relational information and contextual relationships.

Algorithm 1 Knowledge graph-Enhanced Similarity Analysis

Input: Preprocessed legal articles L, consent documents H, knowledge graph $\mathcal{KG}$
Output: Refined similarity matrix $S_{\mathcal{KG}}$
procedure KNOWLEDGE GRAPH INTEGRATION($V_D, \mathcal{KG}$)
    for each vector $v_d \in V_D$ **do**
        Identify entities in d and map to $\mathcal{KG}$
        Extract relationships and neighbors from $\mathcal{KG}$
        Enrich v_d using relational embeddings
    end for
    Return enriched matrix $V_D^{\mathcal{KG}}$
end procedure
procedure REFINED SIMILARITY($V_L^{\mathcal{KG}}, V_H^{\mathcal{KG}}$)
    for each $l_i \in V_L^{\mathcal{KG}}$ **do**
        for each $h_j \in V_H^{\mathcal{KG}}$ **do**
            $S_{ij}^{\mathcal{KG}} \leftarrow \text{CosineSim}(l_i, h_j)$
        end for
    end for
    Return refined similarity matrix $S_{\mathcal{KG}}$
end procedure

c - Annotation

Using the similarity scores, consent document paragraphs are annotated with relevant legal references. The algorithm ensures that only paragraphs with a similarity above a threshold T are annotated, focusing on contextual relevance.

Algorithm 2 Annotation Process

Input: Similarity matrix S, threshold T
Output: Annotations for consent document paragraphs
procedure ANNOTATE PARAGRAPHS(S, T)
    for each paragraph p_i in the consent document **do**
        $S_{i,\max} \leftarrow \max_j S_{ij}$
        if $S_{i,\max} \geq T$ **then**
            Annotate p_i with legal article corresponding to $S_{i,\max}$
        end if
    end for
    Return annotations
end procedure

2.3.3 Experimental results

In this section, we present our experimental setup that is relying on both similarity algorithms parameters, and the use of knowledge graph embeddings.

To begin, we compare the performance of each similarity method by analyzing their ability to capture the semantic alignment between legal articles and consent documents. By applying these methods to our dataset, we generate similarity scores that reflect the degree of resemblance between each document pair, and by using also the knowledge graph to profound the measure of similarity. The distribution of similarity scores obtained from the analysis provides valuable insights into the nature of the relationships between legal articles and consent documents. We present a comprehensive evaluation of different similarity calculation methods and their implications for understanding the relationships between legal articles and consent documents.

We explore as well, the impact of key parameters—max length, window size, and vector size—on the performance of the algorithms in capturing semantic similarity. The algorithms under consideration are Word2Vec, FastText, and XLNet. Each algorithm's architecture and training methodology interact with these parameters differently, affecting their ability to model semantic relationships between words and phrases.

2.3.3.1 Implementation Process

The implementation of the legal compliance component was carried out in a reproducible environment using Google Colab, which provided a GPU-backed infrastructure for running embedding models and managing ontology-based annotations efficiently. The workflow was modularized into four main stages: data preparation, embedding generation, similarity computation, and semantic annotation.

1. Data Preparation. The first stage involved structuring both legal and consent corpora for analysis. Sixty legal articles were extracted from applicable healthcare regulations (including French public health codes and GDPR-related texts), alongside sixty corresponding consent document paragraphs. Each text segment was preprocessed using tokenization, lowercasing, stopword removal, and lemmatization. The textual data was then linked to its corresponding ontology classes (e.g., `LegalObligation`, `RightToWithdraw`, `InformationDuty`) to ensure alignment between the raw corpus and the underlying semantic model.

A preliminary knowledge graph was constructed using RDF triples where each article and consent paragraph was modeled as an individual entity.

2. Embedding Generation. Each text segment was transformed into numerical vectors using pre-trained embedding models (Word2Vec, FastText, and XLNet). The embeddings were generated directly within the Colab notebook, leveraging TensorFlow and PyTorch frameworks for efficient GPU utilization. Parameters such as vector size, context window, and sequence length were dynamically tuned across multiple runs to determine optimal configurations. Each generated embedding was stored in an indexed vector space (FAISS structure) for rapid similarity retrieval.

3. Knowledge Graph Embedding Integration. To bridge symbolic and distributed representations, we integrated the embeddings into the knowledge graph. Using the `pyRDF2Vec` and `RDFlib` libraries, we computed graph embeddings for legal and consent entities, enriching them with contextual features derived from their ontological neighborhood. Object properties such as `partOfRegulation`, `definesRight`, and `appliesToConsentSection` were explicitly modeled to capture legal dependencies. The final representation of each node combined its text embedding and graph-based vector, enabling hybrid semantic similarity evaluation.

4. Similarity Computation and Annotation. Semantic similarity was computed using cosine distance across the hybrid embeddings. The similarity matrices were analyzed to identify the best-matching legal article for each consent paragraph. Matches exceeding a similarity threshold (typically 0.8) were automatically annotated in the knowledge graph through the `isCompliantWith` property, while those below the threshold were flagged for manual review.

5. Validation and Visualization. Finally, similarity distributions and compliance mappings were visualized directly in Colab using `matplotlib` and `networkx`. This allowed inspection of the knowledge graph structure, showing clusters of high-similarity nodes and their legal links. Comparative results confirmed that the hybrid embedding approach (text + KG) outperformed text-only baselines in both precision and interpretability.

2.3.3.2 Experimental Setup

For each algorithm, we conducted a series of experiments varying one parameter at a time while keeping the others constant. This approach allows us to isolate the effects of each parameter on the algorithm's performance in capturing semantic similarity.

We limited our experiments to Word2Vec, FastText and XLNet because they cover three complementary embedding paradigms. Word2Vec provides a robust and computationally efficient baseline for semantic similarity. FastText improves on lexical variation through subword modeling, which is particularly relevant in medical texts. Finally, XLNet introduces contextualized embeddings suitable for capturing long-range dependencies in consent documents. Together these models represent a balanced and widely adopted set of approaches in clinical and legal NLP, enabling fair comparison without incurring the computational overhead of larger transformer architectures.

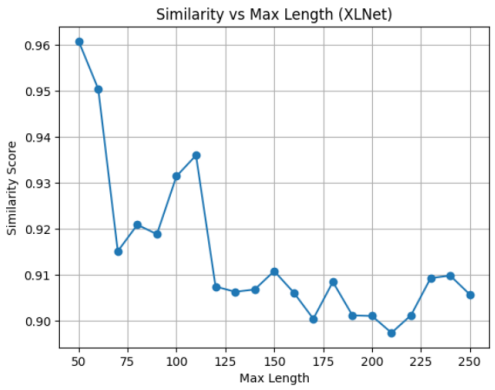
We also embedded the knowledge graph as medium to conduct similarity measurements. The experiments were made on 60 law articles and consent paragraphs. In the following section we showcase

- **Algorithms**: Word2Vec, FastText, XLNet
- **Evaluation Metric**: Cosine similarity between word vectors.

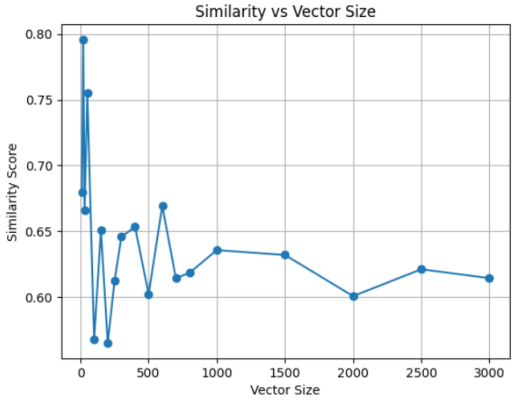
- **Parameters** :
  - Max Length : Maximum number of tokens processed per input sequence, {128, 256, 512, 1024}.
  - Window Size : Context window size around each target word, {2, 5, 10, 15}.
  - Vector Size : Dimensionality of the word vectors, {50, 100, 200, 300}.
- Methodology : Plain similarity measure, Knowledge graph embeddings.

2.3.3.3 Parameter selection fo similarity algorithms

The figure 2.4 shows the variation of cosine similarity in regards to max length and window size for the used similarity algorithms (XLNet, Word2Vec and FastText).



(c) Cosine similarity vs. max length for XLNet



(d) Cosine similarity vs. window size for Word2Vec and FastText

Figure 2.4 – Parameter selection for similarity algorithms

Impact of Parameters on Similarity:
Key parameters like sequence length, context window, and vector size influence the performance of semantic similarity algorithms. For XLNet, a maximum length of 512 balances context and computational efficiency, with marginal gains at 1024. Word2Vec and FastText perform best with moderate context windows (5–10), capturing semantic relationships without introducing noise from irrelevant words. Similarly, vector sizes between 100 and 200 dimensions provide an optimal balance across all models, with larger sizes (300) offering limited improvement at the cost of higher computation. These findings highlight the need to fine-tune parameters to optimize performance and efficiency.

2.3.4 Knowledge graph embedding and annotation

The use of a knowledge graph significantly improves the accuracy of semantic similarity calculations compared to methods without Knowledge graphs. Figure 2.5 illustrates the com-

parison of similarity scores across text-paragraph pairs, showing consistently different scores when leveraging the Knowledge graph. This difference is particularly notable in capturing long-range dependencies and domain-specific nuances, where the Knowledge graph aids in enriching context and resolving ambiguities.

In order to determine whether the use of Knowledge graph allows semantic similarity across the dataset, we compute the difference between the best and second-best similarity matches. The larger differences observed in Knowledge graph-based method suggest a higher confidence in the most relevant match, whereas non-Knowledge graph methods show narrower gaps, indicating less discriminative power. These findings underline the effectiveness of using Knowledge graphs in enhancing both the precision and interpretability of semantic similarity systems.

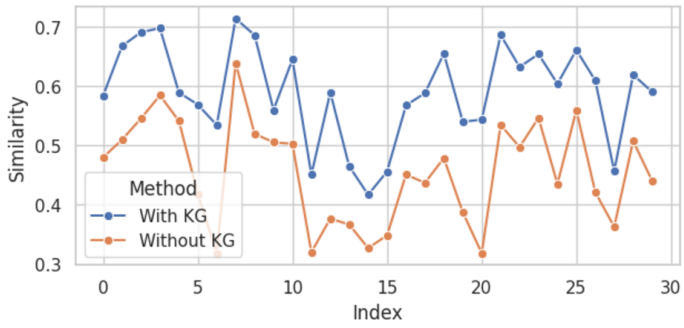


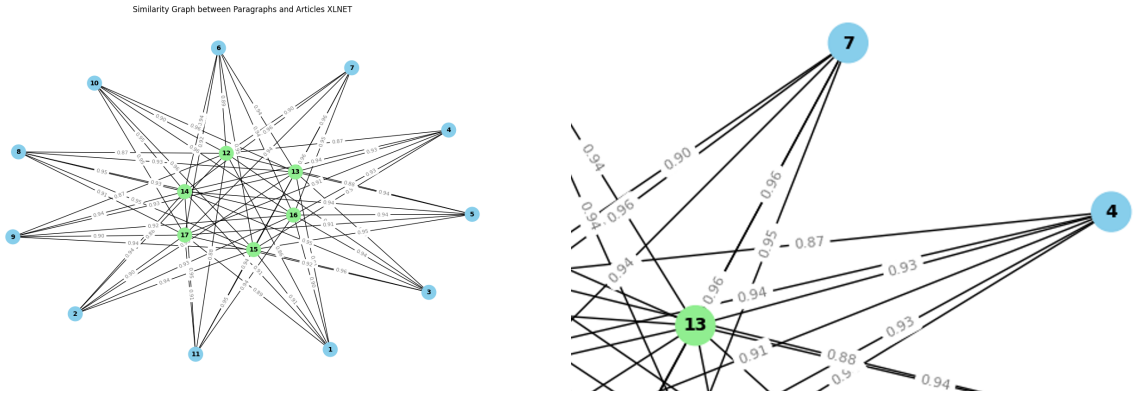
Figure 2.5 – Similarity measures for methods with and without Knowledge graph

The differences in results arise from variations in data quality, hyperparameters, and model architectures. Pre-trained models like BERT and RoBERTa excel in contextual embedding compared to static embeddings from Word2Vec, which depend on parameters like window size and vector dimensions. Models like XLNet and ELMo further leverage bidirectional architectures to capture richer semantic relationships.

The similarity scores play a central role in the annotation mechanism, as they quantify the semantic alignment between consent document paragraphs and legal articles. Figures 2.6a and 2.6b illustrate the results obtained based on XLNet embeddings. The first figure provides a global view of the mapping graph, where each consent paragraph (blue nodes) is connected to relevant legal articles (green nodes), with edge weights proportional to the similarity score. The second figure offers a zoomed-in visualization of the mapping, highlighting the similarity scores.

These results demonstrate that semantic annotation through our approach effectively identifies regulatory correspondences, supporting automated compliance verification and facilitating legal–clinical alignment in consent documentation.

Limits : The proposed approach, while effective, has certain limitations. The performance is highly dependent on the quality and size of the training data, particularly for pre-trained models, which may not fully capture domain-specific nuances. The reliance on similarity scores introduces potential challenges, such as sensitivity to variations in phrasing



(a) XLNET results (b) Similarity measures

Figure 2.6 – Generated similarity annotations

or terminology, which can affect annotation accuracy. Additionally, computational complexity increases significantly with larger models and extended sequence lengths, posing challenges for scalability. Finally, the use of a knowledge graph, while enhancing semantic context, requires careful curation and maintenance, which can be resource-intensive. Addressing these limitations in future work could further improve the system's robustness and scalability.

2.3.5 Synthesis

This section introduced the **Ontology-Driven Machine Learning Annotation for Legal Compliance** approach, designed to ensure that consent documents adhere to the applicable legal standards. By combining natural language processing with ontology-based semantic modeling, the approach enables both automated document analysis and compliance verification.

The proposed approach, while effective, has certain limitations. The performance is highly dependent on the quality and size of the training data, particularly for pre-trained models, which may not fully capture domain-specific nuances. The reliance on similarity scores introduces potential challenges, such as sensitivity to variations in phrasing or terminology, which can affect annotation accuracy. Additionally, computational complexity increases significantly with larger models and extended sequence lengths, posing challenges for scalability. Finally, the use of a knowledge graph, while enhancing semantic context, requires careful curation and maintenance, which can be resource-intensive. Addressing these limitations in future work could further improve the system's robustness and scalability.

The use of the LOINC document ontology proved essential in structuring and aligning the content of legal and consent texts, facilitating similarity consistency. Experimental results demonstrated that integrating ontology-driven features enhances the precision of document similarity annotation. These outcomes confirm the value of adding ontologies on top of statistical methods to advance legal conformity assessment within healthcare documentation.

2.4 Conclusion

This chapter tackled the principal scientific challenges underlying the modeling of informed consent within its legal, ethical, and medical dimensions through the proposed approach SPIRIT — Semantic Personalized and Compliant System for Informed Consent and Care Recommendation. It also addressed the specific challenges related to the first core component of SPIRIT: legal compliance. A major focus was placed on developing an ontology-driven approach capable of ensuring consistent semantic annotation of consent documents.

To meet this objective, we implemented a hybrid methodology adding ontological structuring with textual embeddings. The ontology provided the semantic backbone for representing legal norms and consent concepts, and supported compliance verification. Transformer-based embeddings (XLNet) were employed to capture contextual similarities between legal articles and consent document segments, outperforming traditional static embeddings such as Word2Vec and FastText.

Experimental findings demonstrated that contextual embeddings significantly improve the accuracy of similarity assessments, particularly when integrated with ontology-based representations. These results confirm the effectiveness of adding ontology embeddings with NLP methods to achieve both precision and regulatory conformity in consent document analysis.

Having established a legally sound and semantically coherent foundation, the next chapter 3 builds upon this layer to address the second core component of SPIRIT: personalization. It explores how consent content can be dynamically adapted to individual patient profiles using semantic prompting, user modeling, and generative AI techniques, thereby advancing toward a truly patient-centered and compliant e-consent system.

Chapter 3

Semantic Personalization of Informed Consent

Sommaire

3.1 Introduction . . . 67
3.2 Related Works . . . 69
3.2.1 LLM-Based Personalization Systems . . . 69
3.2.2 Profile ontologies . . . 71
3.2.3 Identified Limitations and Contribution of Our Work . . . 75
3.3 Semantic Personalization over LLMs . . . 75
3.3.1 Formal problem definition . . . 76
3.3.2 Retrieval Layer . . . 77
3.3.3 Semantic Paths Extraction Layer . . . 78
3.3.4 Semantic Reasoning Layer . . . 81
3.3.5 Dynamic Prompt Engineering Layer . . . 84
3.4 Experimental Results . . . 85
3.4.1 Datasets . . . 86
3.4.2 Environment configuration . . . 86
3.4.3 Evaluation metrics . . . 87
3.4.3.1 Statistical Approaches . . . 87
3.4.3.2 Model-Based (Non-LLM) Approaches . . . 88
3.4.3.3 LLM-Evaluation Approaches . . . 88
3.4.4 Evaluation Setup . . . 89
3.4.5 Results . . . 92
3.5 Conclusion . . . 94

3.1 Introduction

Informed consent remains a cornerstone of patient autonomy and ethical medical practice. However, as discussed in Chapter 1, traditional consent processes and documents—often designed to ensure legal compliance—tend to rely on standardized and static information. While

such approaches facilitate regulatory alignment, they frequently fail to account for individual patient characteristics, such as literacy level, cultural background, health condition, or personal preferences. As a result, consent documents may formally satisfy legal requirements while still falling short in terms of patient comprehension and engagement, ultimately weakening the ethical validity of the consent process.

In the previous chapter (Section 2), we introduced SPIRIT, a semantic architecture designed to guarantee the legal and normative compliance of medical consent through ontology-driven reasoning. This first contribution established a robust and transparent framework for structuring consent information according to legal and ethical constraints. However, this compliance-oriented foundation also highlighted a key limitation: legal conformity alone does not guarantee that consent is truly informed. The absence of personalization emerged as a critical gap between formal compliance and meaningful patient understanding.

Addressing this gap requires shifting from generic, one-size-fits-all consent documents toward adaptive and patient-centered information delivery. To ensure that consent is genuinely informed, the content must be tailored to the patient's cognitive, medical, and contextual profile. This transition naturally extends the SPIRIT framework by focusing not only on what information must be provided, but also on how it should be communicated to each individual patient.

Recent advances in Artificial Intelligence, and particularly in Large Language Models (LLMs), offer new opportunities for dynamic content generation and contextual adaptation. LLMs are capable of producing fluent and context-aware text; however, when used in isolation, they lack explicit grounding in medical knowledge, patient models, and personalization logic. As identified in the previous chapter, this limitation poses risks in safety, transparency, and consistency. To overcome these challenges, integrating semantic knowledge—formalized through ontologies and knowledge graphs—becomes essential. Such integration enables structured reasoning over patient profiles, clinical contexts, and consent constraints, ensuring that personalization remains controlled, explainable, and aligned with both legal and ethical requirements.

This chapter presents the second contribution of our work, which builds directly upon the compliance framework introduced earlier and focuses on the semantic personalization of consent documents according to individual patient profiles. Our proposed approach, Semantic Personalization over LLMs, combines the generative capabilities of language models with ontology-based reasoning to deliver consent information that is personalized, understandable, and contextually relevant. First, we review related work on LLM-based personalization and knowledge-grounded text generation (Section 3.2), with particular attention to personalization strategies and modeling choices. Next, we present our proposed approach (Section 3.3), detailing how patient profile ontologies and personalization rules are integrated into LLM prompting. Finally, we describe and discuss the evaluation of the approach (Section 3.4), using both traditional metrics—such as relevance, novelty, and coverage—and emerging evaluation paradigms such as LLM-as-a-judge to assess the impact of semantic personalization.

3.2 Related Works

The notion of informed and voluntary consent, as a legal and ethical principle, ensures that patients are adequately informed about proposed interventions and that their participation remains voluntary and autonomous [108].

However, as medical practices and procedures become increasingly complex, obtaining truly informed consent has become more challenging. Patients often struggle to interpret technical or legal language, while healthcare professionals face difficulties in explaining to individuals with different comprehension levels [108, 109]. These challenges motivate the development of personalized consent systems capable of enhancing understanding, engagement, and ethical aspects.

This section reviews existing research relevant to our contribution on personalized consent generation. It is divided into two parts: first, we review LLM-based personalization approaches (section 3.2.1), which leverage LLMs for adaptive communication; then, we discuss ontology representations (section 3.2.2), which structure and reason over user profiles and medical semantics. Together, these two research directions provide the conceptual and technical foundations of our hybrid ontology–LLM approach.

3.2.1 LLM-Based Personalization Systems

Informed consent remains a key challenge in medical communication, particularly given the technicality of clinical language and the variability of patient literacy levels [110, 111]. The emergence of Large Language Models (LLMs) has opened promising avenues for generating accessible and contextually appropriate medical explanations. Studies show that domain-adapted models trained on medical corpora can produce more understandable and accurate descriptions than traditional expert-authored documents [112]. Applications range from radiology report summarization to generating patient-friendly documentation [113].

Nevertheless, the deployment of LLMs in healthcare raises important limitations. Developing specialized models requires large annotated datasets, domain adaptation, and significant computational resources [114]. More critically, most LLMs exhibit limited reasoning ability and insufficient adaptation to user-specific contexts—capabilities essential for informed consent generation.

To address these issues, several domain-specific variants have emerged, such as BioBERT [115], ClinicalBERT [116], and MedNLI [117]. These models demonstrate improved performance in biomedical entity recognition and inference but remain poorly suited for reasoning about consent or personalizing information to patient attributes. Their fine-tuning processes are data-intensive, static, and difficult to generalize to new personalization factors.

Recent advances have thus turned toward integrating LLMs with structured knowledge representations. Models such as OreoLM [118] and GreaseLM [119] embed knowledge graph

(KG) reasoning within transformer architectures, improving factual grounding and interpretability. Knowledge graphs provide an explainable representation of medical concepts and relationships [86], but most current systems leverage them primarily for question answering or retrieval—rather than adaptive personalization.

Retrieval-Augmented Generation (RAG) architectures have further enhanced factual consistency by introducing external retrieval into the generation process. These systems dynamically fetch relevant entities or documents to enrich model outputs [120, 121]. However, while RAG improves factual coverage and coherence, it lacks explicit reasoning and does not adapt to user-specific profiles—limiting its applicability to individualized consent communication.

In parallel, personalization-oriented approaches have emerged, aiming to tailor generated content to user profiles. Current strategies can be categorized into three main paradigms:

- **Fine-tuning-based personalization**, where models are trained on user-specific data to adapt outputs to known attributes. These approaches achieve deep adaptation but are resource-intensive and inflexible, requiring retraining for each new personalization factor [122].
- **Retrieval-augmented personalization**, which supplements the generative process with profile-relevant data. Although effective for factual enrichment, these models rarely perform reasoning or dynamic adaptation to evolving user states [123].
- **Knowledge graph–enhanced personalization**, which integrates structured reasoning into generation pipelines. Systems such as [124, 125, 126] use KGs to represent user or contextual information and improve coherence and explainability. Yet, they often face challenges in scalability, interpretability, and linguistic fluency.

More recent frameworks, such as [127, 128], propose hybrid graph–text architectures to bridge symbolic reasoning and generative flexibility. Likewise, research on persona-based personalization [129, 130] emphasizes long-term, user-centric adaptation. Despite these advances, most existing methods rely on vector-based representations and remain limited in explainable reasoning.

In summary, current LLM-based personalization systems excel at generating fluent and contextually rich content but lack the semantic grounding required for interpretable and adaptive consent personalization. These limitations motivate the integration of ontologies—explicit, machine-interpretable representations of domain and user knowledge—to enable transparent and rule-based reasoning over personalization decisions.

Synthesis The review of LLM-based systems demonstrates significant progress in automating and simplifying medical communication. However, most approaches remain centered on text generation, summarization, or retrieval, with limited capacity for contextual reasoning or patient-specific adaptation. While hybrid frameworks such as RAG and KG-augmented

LLMs improve factual accuracy and coherence, they largely treat knowledge as external retrieval support rather than an active reasoning layer. Moreover, personalization in these systems is often implicit—driven by statistical similarity rather than an explicit model of patient profiles, cognitive abilities, or health conditions. This indicates that current LLM-based consent and communication systems, although linguistically powerful, still lack the structured personalization and explainability required for ethically compliant informed consent.

3.2.2 Profile ontologies

Ontology-based personalization provides a complementary paradigm to data-driven language modeling by introducing explicit semantics for user attributes, contextual factors, and personalization rules. Unlike purely neural models, which rely on statistical co-occurrences, ontologies enable structured reasoning and explainable decision-making—two essential requirements for ethically compliant consent generation.

Ontologies have been widely used for personalization across multiple domains, such as e-learning, cultural heritage, and e-commerce [131, 132]. In healthcare, they have supported adaptive clinical decision systems, patient-centered recommendation engines, and context-aware monitoring [133, 134]. These approaches typically rely on formal knowledge structures that represent user characteristics, environmental context, and system interactions.

Several ontologies have been proposed specifically for modeling user or patient profiles:

- The **Friend of a Friend (FOAF)** ontology is one of the earliest attempts to describe individuals, their attributes, and relationships in a machine-readable format. Although generic, it provides a foundational schema for user modeling [135].
- The **Health Level 7 (HL7)** and the **SNOMED CT** ontology focus on clinical and administrative data interoperability, representing medical history and healthcare processes [136, 7].
- **Enhanced Personal Profile Ontology (E-PPO)** : provides a domain-oriented schema that captures patient demographics, health status, in the heterogeneous domain of complex and Chronic health Conditions, and focuses on modeling personal characteristics and preferences for adaptive systems [137].
- **Clinical Smell and Taste Disorders Ontology (CSTDO)** : Models demographic factors, patient descriptors, and symptomatology in a structured way, making it suitable for representing general patient profiles beyond its original domain [138].
- **Health and Social Person-Centric Ontology (HSPO)** : Designed to support person-centered care, HSPO integrates clinical health data with social, environmental, and contextual determinants that influence health outcomes [139].
- **Culture, Race, and Ethnicity Ontology (CRENO)** : It provides a semantic framework for representing cultural, racial, and ethnic identity dimensions relevant to healthcare delivery and health equity research [140].

While these ontologies offer structured representations of patients or users, they tend to either (i) focus on demographic and relational information without linking it to experiential or perceptual dimensions, or (ii) concentrate on clinical and procedural information without modeling personal or contextual variability. For our objective of personalizing informed consent, both aspects are essential: demographic and medical attributes must coexist with representations of perception, comprehension, and patient experience.

To address this gap, our work leverages the **Clinical Smell and Taste Disorders Ontology (CSTDO)** [138], which was initially developed to model sensory dysfunctions but includes rich conceptual structures describing:

- **Demographic attributes** (e.g., age, gender, nationality);
- **Medical and pathological characteristics** (e.g., symptoms, conditions, severity levels);

These features make CSTDO particularly suitable for modeling personalized consent profiles. In our dataset, patient profiles include comparable information — age, literacy, pathology, and linguistic or cognitive indicators — which align closely with the demographic and perceptual constructs defined in CSTDO. This semantic compatibility facilitates the mapping between real-world data and ontology classes, ensuring a consistent and interpretable representation of patient characteristics.

Moreover, the integration of CSTDO with the **LOINC Document Ontology** (introduced in Chapter 1) enables a comprehensive link between document semantics and patient profiles. LOINC captures the typology and medical context of consent documents (e.g., "type of service," "kind of document"), while STDO structures user-specific factors that influence how information should be adapted or communicated. Together, these ontologies provide a dual-layer semantic foundation for reasoning about both the content to be personalized and the profile guiding personalization.

In summary, ontology-based personalization allows for transparent and explainable adaptation mechanisms that align with both legal and ethical dimensions of informed consent. By using CSTDO, we exploit a knowledge base that not only encodes medical and perceptual dimensions but also aligns naturally with the real demographic and personal attributes found in our dataset. This alignment ensures that the personalization reasoning layer operates on semantically grounded, interoperable data representations.

Synthesis Ontology-driven personalization approaches provide the semantic and logical rigor missing in purely data-driven models. By explicitly modeling user characteristics, contextual parameters, and reasoning rules, they enable interpretable adaptation mechanisms aligned with ethical and regulatory requirements. Nonetheless, most existing ontologies either overemphasize clinical descriptors (e.g., SNOMED CT, HL7) or remain too generic to capture perceptual and demographic diversity (e.g., FOAF, E-PPO). The Clinical Smell and Taste Disorders Ontology (CSTDO) bridges this gap by including demographic, perceptual,

and experiential concepts that align closely with real patient data used in our work. Its integration with the LOINC Document Ontology provides a coherent dual-layer model linking consent document semantics and patient profile attributes. Together, these ontologies offer a robust foundation for reasoning about how consent content should be adapted to individual profiles, supporting transparent and explainable personalization.

Paper	Dataset	Retrieval Technique	Type of LLM	Results	Metrics
KG-RAG [124]	CWQ (complex questions, SPARQL queries)	Vector similarity search (Sentence-Transformer + RedisDB)	GPT-4 Turbo (Preview)	RAG Benchmarks: QA: Exact Match (EM), F1 Score, Hallucination rate	EM, F1, Hallucination rate
BioRAG [125]	23M PubMed abstracts, external DBs (Gene DB, dbSNP)	MeSH-based probabilistic model + vector similarity search	Llama3-8B (fine-tuned), Llama3-70B (base), Pubmed-BERT	Fine-tuning on BioLLMs, Benchmarks: PubMed	F1, Precision, Recall
Graph-Reader [127]	HotpotQA, 2WikiMulti-hopQA, MuSiQue, NarrativeQA	BM25 (top-1, top-3), Ada-002, GPT-4-based retrieval	GPT-4-128k	Comparison of retrieval techniques (LR-1, LR-2)	QA: F1, EM, LLM Rating
MSTS5 [141]	LC-QuAD 2.0, QALD-9, QALD-10	Fine-tuning	mT5 pretrained on mC4 (fine-tuning 300 epochs, early stopping)	Comparison with other KGQA models	F1, Precision, Recall
Chain-of-Thought [11]	-	Few-shot examples	LaMDA, GPT-3, PaLM, UL2-20B	-	-
Chain-of-Knowledge [142]	WikidataSm, KnowReason	Knowledge integration	-	-	-
Two-stage Generative QA on Temporal KGs [143]	GTEA (temporal KGQA tasks)	Subgraph retrieval (time-based ranking, time mining)	Open-source instruction-tuned LLMs	Novel ICL framework for TKGQA	EM, F1
Knowledge Graph-Enhanced LLMs via Path Selection [126]	FreebaseQA, WikidataSm	Path-based retrieval	Fine-tuned GPT models	RAG with KG paths improves reasoning	F1, Precision
Learning to Plan for RAG-Large LLMs from KGs [128]	SQuAD, HotpotQA	Planning-based retrieval	Llama2-13b	KG-based task planning for RAG	Enhanced QA-specific metrics
DALK: Dynamic Co-Augmentation for Alzheimer's QA [144]	AD-related corpus (PubMed), Alzheimer's Ontology	KG-driven augmentation	GPT-4, fine-tuned BioLLMs	Outperforms baselines in Alzheimer's QA	EM, F1, specificity, sensitivity

Table 3.1 – Comparison of Knowledge Graph-Augmented Retrieval Techniques using LLMs.

3.2.3 Identified Limitations and Contribution of Our Work

Across the reviewed literature, three major limitations persist:

1. **Limited reasoning capabilities.** Most existing LLM-based consent systems focus on retrieval and summarization rather than structured inference or semantic reasoning.
2. **Weak personalization grounding.** Personalization is often heuristic, relying on embeddings or retrieved examples, without a formal representation of user profiles or adaptation rules.
3. **Lack of interpretability.** Current methods rarely provide traceable, explainable justifications for personalized output—an essential feature for ethical and legal transparency in healthcare.

Our proposed framework directly addresses these gaps by integrating **ontology-based reasoning** with **LLM-driven personalization**. Specifically, we:

- Formalize patient profiles and consent-related knowledge using ontologies and knowledge graphs;
- Derive explicit personalization decisions through logical inference;
- Use these decisions to condition LLM generation via structured prompts, ensuring personalized, explainable, and legally aligned consent generation.

3.3 Semantic Personalization over LLMs

Building on the legal component established in the previous chapter, this section introduces the second component of the SPIRIT architecture: **Semantic Personalization**. Its objective is to adapt consent content to individual patient profiles through the hybrid use of ontological reasoning and LLMs.

While traditional personalization methods rely mainly on embeddings or template-based adjustments, our approach leverages the structured semantics of knowledge graphs and the generative flexibility of LLMs. By aligning the consent knowledge graph with patient profile ontologies, semantic reasoning enables the extraction of explicit personalization rules, which are then dynamically integrated into LLM prompts. This design ensures that the generated consent text is both contextually tailored and semantically traceable.

The proposed architecture, illustrated in Figure 3.1, is composed of four interconnected layers: (1) Retrieval Layer, (2) Paths Extraction Layer, (3) Semantic Reasoning Layer, and (4) Dynamic Prompt Engineering Layer. Throughout this section, we start by giving a formal problem definition of the personalization use case (in section 3.3.1). Then, we detail each of the architecture's layers.

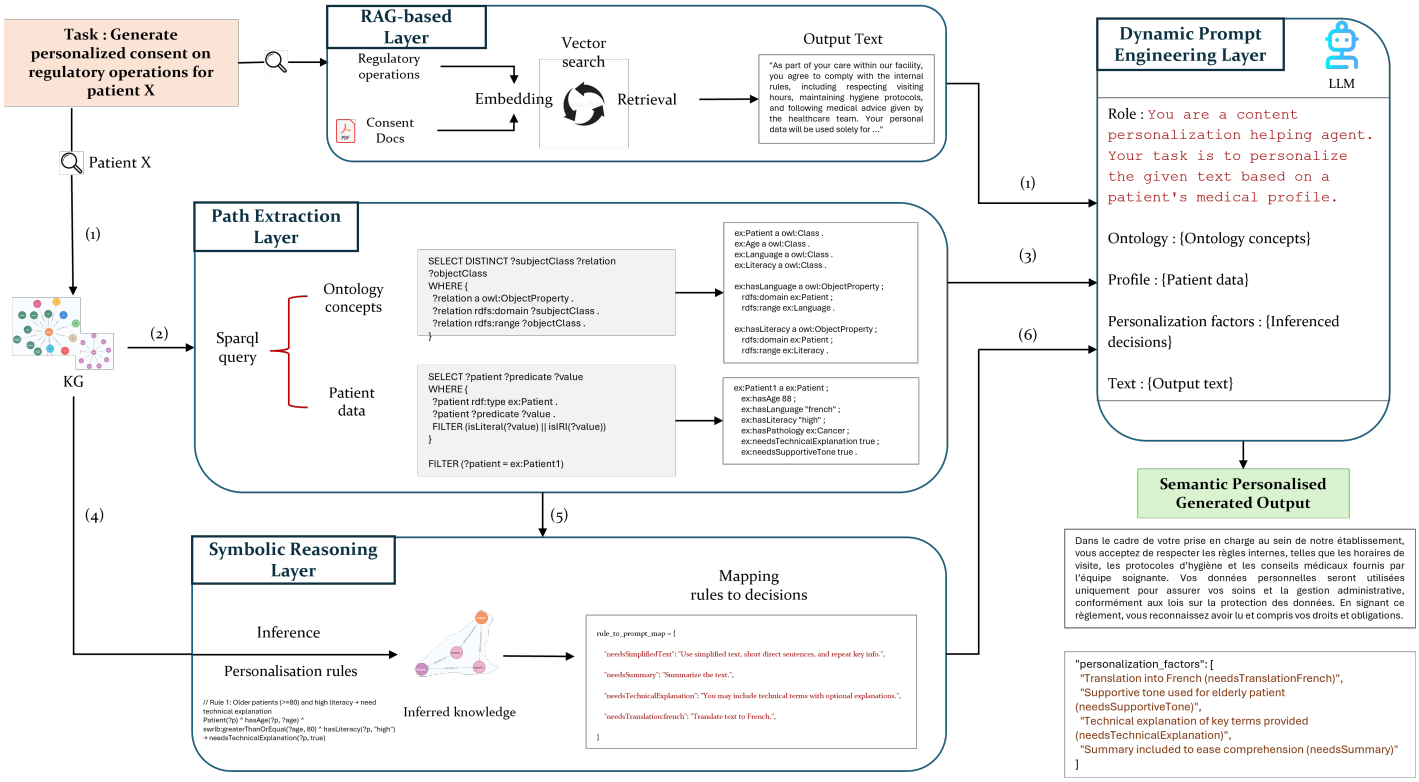


Figure 3.1 – Semantic Personalization over LLMs approach

3.3.1 Formal problem definition

Personalizing informed consent requires more than generating alternative text—it involves adapting medical, ethical, and contextual information to the specific needs of a patient while ensuring accuracy and compliance. To clearly frame this challenge, we formalize the problem as follows:

Let P be a patient profile, encoded as RDF triples. It may include medical history, demographic attributes, behavioral data, personal preferences, and any contextual information relevant to consent. Let T be a generic segment of a consent document, such as a paragraph describing a procedure, risks, or benefits. Let O be a domain ontology that models the medical concepts, relationships, and semantic constraints associated with the consent process (e.g., diseases, treatments, risks, contraindications). The objective is to generate a personalized version T' of T, such that:

- T' remains faithful to the original medical and legal meaning of T,
- T' is adapted to the patient's characteristics described in P (e.g., language level, medical condition, comprehension needs),

Let: $P \subseteq \mathcal{G}, \quad T \in \mathcal{L}, \quad O \in \mathcal{K}$ where:

- $\mathcal{G}$ is the space of RDF graphs representing patient profiles;
- $\mathcal{L}$ is the space of linguistic segments of consent documents;
- $\mathcal{K}$ is the space of domain ontologies modeling medical and legal concepts.

The personalization process is decomposed into two stages:

(1) Semantic Reasoning. A decision function $d : (\mathcal{G}, \mathcal{K}) \to D$ computes a set of personalization decisions D from the combination of a patient profile P and ontology O:

$$D = d(P, O) = \text{Reason}(P \cup O)$$

where Reason denotes inference over the merged knowledge graph (e.g., via SWRL rules).

(2) Personalized Text Generation. A generation function $f : (\mathcal{L}, \mathcal{K}, D) \to \mathcal{L}$ produces a personalized text T' conditioned on the original text T, ontology O, and personalization decisions D:

$$T' = f(T, O, D) = \arg\max_{T'} p(T' \mid T, O, D; \theta_{LLM})$$

where θ_{LLM} are the parameters of the language model used for generation.

This formulation expresses the personalization pipeline as a composition of reasoning and generation:

$$T' = (f \circ d)(T, O, P)$$

enabling interpretable, ontology-guided personalization of consent information.

3.3.2 Retrieval Layer

The personalization process begins with the identification of relevant textual content within the consent document. The **Retrieval Layer** serves as the entry point of the architecture, responsible for extracting semantically relevant document segments that will later be adapted to the patient profile.

To achieve this, we employ a Retrieval-Augmented Generation (RAG) mechanism that queries the knowledge base to locate consent clauses, norms, or document fragments related to the personalization task. This retrieval ensures that subsequent reasoning and generation steps operate on targeted and contextually coherent content, rather than on the entire document.

In our case, since personalization can concern various aspects of a consent document, the retrieval layer isolates the corresponding sections or "chunks" most relevant to each personalization query. The main steps of this process are detailed below.

- Task: The task is the input query that sets context to the components of the approach. It contains several information, first on the action that the LLM should do (in terms of content generation). And second, information related to the consent forms, on which information to retrieve.
- Embedding: Using embedding models, like BERT, we encode and condense passages of text from the corpus as dense vector representations. An incoming user prompt is encoded as a vector representation, using the same embedding model.
- Retrieval: Using the same embedding model, documents are embedded into a vector base. Using distance metrics like cosine similarity, the system runs a search over the encoded passages to find closest neighbors to the query vector.

The output of this component is the text (or list of texts) related to the task, and that need personalization, which will be the aim of the next layers.

3.3.3 Semantic Paths Extraction Layer

Once the relevant textual segments are retrieved, the next step is to identify the semantic connections that link patient characteristics to document concepts. The Semantic Paths Extraction Layer bridges the profile and document spaces by uncovering relationships between ontology entities—such as patient conditions, literacy level, or language preference—and the consent document's thematic elements.

This layer leverages the structured knowledge encoded in the underlying ontologies, particularly the Clinical Smell and Taste Disorders Ontology (CSTDO) for profile representation. Through SPARQL queries, semantic paths are extracted from the knowledge graph, capturing how specific patient traits interact with consent document concepts.

These paths form the semantic backbone that will guide rule-based reasoning in the subsequent layer.

CSTDO is a domain-specific knowledge resource designed to formally represent and standardize concepts related to the clinical assessment, diagnosis, and classification of olfactory (smell) and gustatory (taste) impairments. The strength of the CSTDO lies in its capacity to define the clinical profile of a patient. Figure 3.2 shows the different used concepts of the ontology.

Beyond the conceptual classes representing disorders, symptoms, causes, and diagnostic methods, the ontology integrates a set of **object properties** and **data properties** that structure relationships between clinical entities and patient instances. Object properties such

as `hasSymptom`, `hasDiagnosis`, `hasCause`, `isAssessedBy`, and `receivesTreatment` establish semantic links between a patient and the corresponding medical concepts. Data properties, on the other hand, describe literal attributes such as `age`, `gender`, `literacyLevel`, and `languagePreference`, which are essential for personalization. Together, these relations enable a fine-grained, machine-interpretable representation of a patient's medical and contextual profile.

In our implementation, the patient profile is represented as an **individual** in the ontology, connected to multiple medical and demographic concepts through these relations. For instance, a patient individual may be linked to a pathology instance (e.g., `Hyposmia`) through `hasDiagnosis`, associated with a therapeutic approach (e.g., `OlfactoryTraining`) via `receivesTreatment`, and described by personal attributes such as `age = 78` or `literacyLevel = low`. These structured relationships form a semantic network that captures both clinical and social determinants of personalization.

The ontology also supports hierarchical reasoning through **class axioms** and **restrictions**, allowing inferred knowledge to emerge dynamically. For example, the property chain `hasDiagnosis → isAssociatedWith → CarePlan` enables automatic inference of appropriate care recommendations, while `hasSymptom some OlfactoryDisorder` can classify a patient as belonging to the broader category of `SensoryImpairment`. Such reasoning capabilities are crucial for generating semantically coherent and medically appropriate personalization factors used in the LLM prompts.

Finally, the CSTDO ontology is aligned with higher-level standards such as **LOINC** and **SNOMED CT**, ensuring interoperability and consistency across healthcare datasets. This alignment supports the integration of external medical terminologies, allowing our semantic framework to generalize beyond olfactory disorders and be extended to broader clinical contexts in future applications.

The mentioned ontology has been chosen because of it's related profile concepts to the actual existing data, and we added other semantic personalization factors to it (language, literacy, etc). In order to achieve the reasoning and prompt engineering afterwards, we need to extract relevant information from the KG. Here are the steps of the process :

Ontology Querying. The ontology provides a formal schema of concepts such as `Age`, `Language`, `Literacy`, and their relationships to the class `Patient`. We query the ontology to extract all relevant object properties and their domain-range constraints. This structural information is used to identify which patient attributes can influence personalization, and ensures extensibility by supporting dynamic schema introspection.

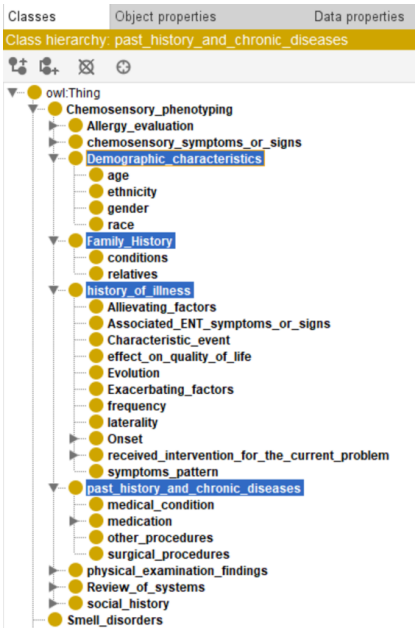


Figure 3.2 – Concepts of CSTDO visualized in Protégé

SPARQL query for extracting ontology structure:

```
SELECT DISTINCT ?subjectClass ?relation ?objectClass
WHERE {
  ?relation a owl:ObjectProperty .
  ?relation rdfs:domain ?subjectClass .
  ?relation rdfs:range ?objectClass .
}
```

Patient Data Querying. The patient's profile is represented as RDF triples, describing both demographic and clinical characteristics (e.g., age, language preference, literacy level, pathology). These values are extracted through a simple pattern-matching query and passed to the symbolic reasoning engine.

SPARQL query for extracting patient-specific data:

```
SELECT ?predicate ?value
WHERE {
  ex:Patient1 rdf:type ex:Patient .
  ex:Patient1 ?predicate ?value .
  FILTER (isLiteral(?value) || isIRI(?value))
}
```

The output of this layer is the set of concepts and Patient data that can be inserted later in the other layers. On one hand, the ontology concepts and the patient data are given to the Dynamic Prompt engineering Layer, which fills the prompt of the LLM, and gets updated for different patient profile. And on the other hand, the output is used as additional

information for the Symbolic Reasoning Layer, that helps in mapping inferred knowledge with the personalization decisions.

3.3.4 Semantic Reasoning Layer

The Semantic Reasoning Layer operates the personalization logic by interpreting the semantic paths obtained earlier. It applies rule-based reasoning to infer personalization needs from the knowledge graph—for example, determining that a patient with low literacy requires simplified text, or that a non-native speaker needs translated consent materials.

This inference is performed using the Semantic Web Rule Language (SWRL), which enables the formal specification of adaptation rules over the ontology. Each rule encodes a condition-action pair linking patient attributes to personalization strategies. The reasoning process enriches the knowledge graph with inferred statements, producing explicit personalization directives that will be used in the next stage.

This layer thus acts as the system's interpretive core, transforming static profile data into actionable adaptation knowledge.

SWRL-Based Personalization Rules The extracted data from the KG is evaluated against a set of SWRL rules, which serve as decision-making logic for personalization needs. These rules capture common semantic patterns such as: "patients aged over 85 require a supportive tone", or "patients with low literacy benefit from simplified content".

SWRL Rules Example :

```
Patient(?p) ^ hasAge(?p, ?age) ^ swrlb:greaterThan(?age, 85)
   -> needsSupportiveTone(?p)

Patient(?p) ^ hasLanguage(?p, "french")
   -> needsTranslationFrench(?p)

Patient(?p) ^ hasLiteracy(?p, "low")
   -> needsSimplifiedText(?p)

Patient(?p) ^ hasPathology(?p, ex:Cancer)
   -> needsTechnicalExplanation(?p)
```

These rules produce inferred personalization decisions as logical assertions (e.g.,needsTranslationFrench), which are then aggregated for subsequent generation steps.

Mapping Inferred Knowledge to Generation Instructions The output of the reasoning step is a set of high-level personalization needs. These are not yet actionable for the

LLM. Therefore, we define a *rule-to-prompt map* that transforms each symbolic inference into a natural language instruction that can guide the LLM's behavior.

Rule-to-Prompt Mapping :

```
{
  "needsSimplifiedText":
    "Use simplified text, short direct sentences, and repeat key info.",

  "needsSummary":
    "Summarize the text.",

  "needsTechnicalExplanation":
    "Include technical terms with optional lay explanations.",

  "needsTranslationFrench":
    "Translate text to French.",

  "needsSupportiveTone":
    "Use a supportive and reassuring tone appropriate for elderly care."
}
```

The Symbolic Reasoning Layer has as output the decisions for personalization (that are deduced based on the mapping with the KG SWRL rules)

The personalization instructions are dynamically composed for each patient, based on the set of SWRL-inferred personalization needs. These are embedded in a prompt template designed to condition the LLM:

Prompt Template:

```
Role: You are a content personalization helping agent.
Your task is to personalize the given text based on a patient's medical
    profile.
Ontology: {Ontology concepts}
Profile: {Patient data}
Personalization factors: {Inferred decisions as instructions}
Text: {input text to be personalized}
```

This prompt template represents the semantic interface between symbolic reasoning and LLM generation. Each component is dynamically built from the knowledge graph and its reasoning output, and serves a distinct function:

- **Role:** Defines the behavior of the LLM, constraining it to act as a personalization assistant rather than a general text generator.
- **Ontology:** Provides the vocabulary and concepts, ensuring that the generated content remains aligned with domain-specific semantics.

- **Profile:** Encodes patient-specific information (e.g., age, pathology, literacy level, language preference).
- **Personalization Factors:** Represents the explicit reasoning output derived from SWRL rules and ontology inference.
- **Text:** The input text that needs to be personalized and adapted (e.g., a consent document).

To refine this process, we initially experimented with various prompt engineering techniques using large language models such as GPT-4. However, these conventional approaches—based solely on static or template-based prompting—did not yield satisfactory personalization results, particularly when adapting consent content to diverse patient profiles. Consequently, we developed a **semantic dynamic prompting** approach, in which prompt composition is guided by ontology reasoning and SWRL-inferred rules. This was combined with a context-grounded conditioning mechanism, enabling the LLM to dynamically adapt to each patient's semantic profile while maintaining domain coherence and legal compliance.

Beyond the prompt structure itself, a key innovation of this layer lies in the use of a **context file** that acts as a unified, dynamic representation of all personalization elements. Instead of embedding static information directly into the prompt, the system maintains a structured file—typically in JSON format—where the ontology, inferred personalization rules, and patient profile data are stored together. This file is automatically generated and updated before each personalization task, ensuring that every prompt is grounded in the most recent and relevant information.

The context file contains:

- The active subset of ontology concepts and relations relevant to the patient and consent type;
- The output of SWRL reasoning, including inferred personalization rules (e.g., "simplify text for low literacy" or "translate to patient's language");
- Patient-specific metadata such as demographics, pathology, language, and comprehension level;
- Compliance constraints and document-specific context derived from the legal layer.

Before each prompt is sent to the LLM, this context file is **registered dynamically** as part of the prompt session. In practice, this means the LLM loads the file at runtime, allowing it to operate as if it were interacting with an evolving internal model of the patient and consent structure. When a new patient is processed or when new reasoning inferences are produced, the file is automatically refreshed and re-registered, ensuring continuous adaptation.

This design turns the prompt into a dynamic semantic bridge between the symbolic and generative layers. Rather than sending a fixed instruction to the LLM, each personalization query is dynamically reconstructed from the current context, providing:

- **Adaptivity:** Each patient's context produces a unique and updated prompt.
- **Traceability:** All personalization decisions can be traced back to ontology and rule-level reasoning.
- **Semantic consistency:** The generated text always remains coherent with the ontology and legal model.

The key advantage of this structure is that the prompt is both **dynamically generated** and **semantically grounded** in symbolic reasoning. Each element of the patient profile contributes to a tailored instruction set, ensuring that the LLM produces consent text that is compliant, medically appropriate, and personally accessible.

3.3.5 Dynamic Prompt Engineering Layer

The final component, the **Dynamic Prompt Engineering Layer**, transforms the inferred personalization directives into natural language instructions for LLM. It serves as the interface between the symbolic and generative components of the system, ensuring that reasoning outcomes effectively guide text generation.

Through a rule-to-prompt mapping strategy, the system converts ontological inferences (e.g., `needsSimplifiedText`, `needsTranslation`) into structured prompt elements integrated within the LLM input. This dynamic prompt construction allows the LLM to adapt its generation process according to the inferred context, ensuring that the resulting consent text is both legally compliant and personally relevant.

This layer therefore closes the semantic-to-generative loop, enabling personalized consent generation grounded in explicit knowledge reasoning.

The dynamically constructed prompt integrates both the **file context** (e.g., the consent document, the legal norms, and their summaries) and the **patient profile context** (e.g., age, language, literacy, pathology). This fusion enables the model to generate a *personalized, compliant, and contextually accurate* version of the consent text.

More precisely, each input document is represented as a set of triples describing its type, associated norms, and linked summaries, while the patient profile provides demographic and clinical attributes. These two contexts are jointly serialized and encoded as a prompt template following the structure:

```
[Document Context]
Norm: <norm_description>
Summary: <summary_text>

[Patient Context]
Age: <age>
Language: <language>
Literacy: <literacy_level>
```

```
Pathology: <pathology>

[Task]
Generate a personalized consent text based on the patient's characteristics
while ensuring clarity, compliance, and relevance. Explicitly list the
personalization factors and the reasoning behind each adaptation.
```

To illustrate, a patient profile such as:

```
ex:Patient1 a ex:Patient ;
  hasAge 88 ;
  hasLanguage "french" ;
  hasLiteracy "low" ;
  hasPathology ex:Cancer .
```

would trigger the following inferred personalization factors (using ontology-driven reasoning rules):

```
"personalization_factors": [
  "Translation into French (needsTranslationFrench)",
  "Supportive tone used for elderly patient (needsSupportiveTone)",
  "Technical explanation of key terms provided (needsTechnicalExplanation)",
  "Summary included to ease comprehension (needsSummary)"
]
```

These personalization factors are automatically converted into explicit **prompt instructions** and injected into the generative pipeline, forming the *semantic prompt* that guides the LLM. The final output includes:

- A **personalized consent text**, linguistically and semantically adapted to the patient.
- The list of applied **personalization factors**, ensuring explainability and traceability.
- The LLM's **reasoning**, describing how each factor influenced the generation.

This mechanism enables the LLM to act as a controlled generative component within a semantically grounded framework, ensuring that the produced text is not only human-readable but also *compliant, explainable, and auditable*. By leveraging structured knowledge and contextualized prompt engineering, the system bridges the gap between symbolic reasoning and generative adaptability.

3.4 Experimental Results

This section presents the experimental evaluation of our proposed semantic personalization approach. The goal of this evaluation is to assess the effectiveness of integrating semantic reasoning with LLMs for generating personalized consent documents. In particular, we aim to quantify how semantic reasoning contributes to improving the relevance, coverage, and personalization quality of generated texts compared to standard LLM-based approaches.

To achieve this, we designed a set of experiments that evaluate the system from both a technical and user-centered perspective. The evaluation examines (i) the influence of ontology-based reasoning on personalization accuracy, (ii) the comparative performance of our approach with and without semantic enrichment, and (iii) the interpretability and consistency of personalization outcomes.

The following subsections detail the experimental framework: the datasets used for document and profile representation (in section 3.4.1), the environment configuration (in section 3.4.2), the evaluation metrics adopted to measure personalization performance (in section 3.4.3), the setup of comparative experiments, and finally, the obtained results with their analysis (in section 3.4.5).

3.4.1 Datasets

Medical Content We use the **WebMD**[1] dataset that is used in the Persona-healthcare work [129], which contains professionally written medical articles and summaries on a wide range of pathologies, treatments, and care instructions, that are in a Json format, and categorized by disease. For this study, we selected a subset of content related to chronic and age-related diseases.

Patient Profiles We simulate a set of synthetic patient profiles, generated to reflect the distribution of real patients from a local healthcare institution. Each profile includes attributes such as age, language, literacy level, and pathology, structured according to our ontology. This allows us to rigorously test personalization strategies under diverse patient conditions.

3.4.2 Environment configuration

The different experiments are run using google Colab GPU T4, with two different LLMs with different temperatures. Llama3.1 for temperatures ranging from 0.3 to 0.7, and Deepseek-r1 with the same range of temperature. Both LLMs are used locally, with Ollama setup.

As for the datasets, each summary from the webmd dataset is personalized for each patient (50 summaries and 10 patients). Which allows us to generate a list of dataframes (of 500 rows-generations for two aspects : with and without ontology, so 1000 generations in total) covering all summaries for all patients, with an output containing three features :

personalized_text	**personalization_factors**	**reasoning**

[1]Webmd persona dataset : https://github.com/ankan2/persona-healthcare

3.4.3 Evaluation metrics

In terms of metric evaluation, we focus on the proposed metrics in literature that we describe in this section. Evaluating Large Language Model outputs is more complex than typical NLP tasks. In essence, a good evaluation ideally satisfy these attributes :

- Quantitative : presenting a numeric score, enabling comparison, thresholds, or continuous improvement tracking.
- Reliable / Stable — scores should not wildly fluctuate under minor variations in prompt or output.
- Accurate — the score should reflect human judgment: relevance, correctness, coherence, faithfulness, etc.

To satisfy the mentioned attributes, We can classify evaluation approaches into three classes : Statistical approaches, Model-based (non-LLM) approaches and LLM-Evaluation approaches

3.4.3.1 Statistical Approaches

Statistical approaches rely on surface-level comparison between generated text and a reference output. They operate on symbolic or lexical overlap, measuring similarity through string- or token-based statistics. Typical metrics include n-gram precision, recall, and overlap ratios, such as BLEU, ROUGE, METEOR, and edit distance.

BLEU (Bilingual Evaluation Understudy) BLEU is one of the earliest and most widely used metrics for evaluating text generation tasks such as translation or summarization. It measures the proportion of n-grams in the generated text that also appear in the reference, balancing precision with a brevity penalty to discourage short outputs. While effective for system benchmarking, BLEU is insensitive to paraphrasing and semantic equivalence.

ROUGE (Recall-Oriented Understudy for Gisting Evaluation) ROUGE focuses on recall rather than precision, measuring how much of the reference text is covered by the generated text. Variants like ROUGE-n (based on n-gram recall), ROUGE-L (based on longest common subsequence), and ROUGE-S (based on skip-bigram co-occurrence) are commonly used in summarization tasks. ROUGE performs well when reference and candidate texts share similar phrasing but struggles with semantically equivalent paraphrases.

METEOR (Metric for Evaluation of Translation with Explicit ORdering) METEOR extends BLEU by incorporating stemming, synonymy, and paraphrase matching, al-

lowing for more flexible word alignment between candidate and reference. It combines precision, recall, and a penalty for fragmented word order. METEOR correlates better with human judgment than BLEU or ROUGE, but it still largely depends on word-level similarity.

3.4.3.2 Model-Based (Non-LLM) Approaches

Model-based approaches introduce semantic awareness by employing pre-trained neural encoders, typically derived from contextual language representations such as BERT or RoBERTa. Instead of comparing surface tokens, these methods embed candidate and reference texts into a shared latent space and compute their similarity using vector-based metrics.

BERTScore BERTScore computes similarity between contextual embeddings of words or tokens in the candidate and reference texts. By leveraging transformer-based encoders, it measures semantic overlap even when the exact words differ. This results in evaluations that are more aligned with human judgments, particularly in tasks involving rephrasing or abstract summarization.

BLEURT BLEURT builds on top of BERT representations but is fine-tuned on human-labeled data, allowing it to directly predict quality scores. Unlike heuristic metrics, BLEURT learns to assess fluency, factual correctness, and coherence, making it adaptive to different domains and tasks. However, its performance depends on the diversity and quality of the data used for fine-tuning.

NLI-Based Evaluation Natural Language Inference (NLI)–based metrics evaluate generated text through entailment reasoning. These models assess whether the generated output logically supports, contradicts, or is neutral with respect to the reference. This approach captures logical consistency and factual correctness, but remains limited by the robustness of the underlying entailment model and its ability to handle complex, domain-specific reasoning.

3.4.3.3 LLM-Evaluation Approaches

The most recent paradigm, known as *LLM-as-a-judge*, leverages large language models themselves to perform evaluation. Instead of relying on explicit similarity measures, an evaluator LLM is prompted to assess a text according to predefined rubrics or quality dimensions, such as faithfulness, coherence, factuality, or ethical alignment.

G-Eval G-Eval introduces structured, rubric-guided evaluation using LLMs like GPT-4. It integrates reasoning steps, or "chain-of-thought" prompts, to produce consistent and interpretable judgments. This approach allows for multi-dimensional assessment (e.g., relevance, fluency, factual accuracy) while aligning more closely with human evaluators.

Prometheus Prometheus represents a formalized version of the LLM-as-a-judge paradigm. It decomposes evaluation into distinct criteria—such as factuality, reasoning quality, and ethical soundness—and prompts the LLM to rate each dimension separately before aggregating the results. This structure reduces prompt bias and enhances reproducibility.

Rubric- and Graph-Based Evaluators Recent approaches use directed or rubric-based frameworks where evaluation criteria are represented as interdependent components (e.g., factuality influencing coherence). LLMs reason through these structures to provide holistic judgments. Such frameworks aim to replicate the depth and interpretability of expert human evaluation while maintaining scalability.

Synthesis Statistical and model-based approaches offer objective, reproducible evaluation grounded in quantitative similarity, but they fail to capture higher-order reasoning, contextual nuances, or pragmatic coherence. LLM-based evaluators bridge this gap by embedding reasoning and comprehension directly into the evaluation process. However, they introduce new concerns—such as sensitivity to prompt design, potential evaluator bias, and limited reproducibility—that continue to motivate hybrid and standardized evaluation frameworks.

In our study, we position the evaluation of our framework within two complementary paradigms: *statistical evaluation* and *LLM-based evaluation*. The statistical evaluation provides quantitative evidence of performance by measuring lexical and structural similarity between the generated and reference texts, ensuring reproducibility and objectivity. In parallel, the LLM-based evaluation leverages the interpretive and reasoning capabilities of large language models to assess the semantic quality, coherence, and contextual appropriateness of the personalized outputs. This dual evaluation strategy enables a balanced assessment—combining traditional quantitative metrics with qualitative, human-aligned judgment—while placing a particular emphasis on the degree of personalization achieved in the generated consent documents.

3.4.4 Evaluation Setup

For each patient-profile and content pair, we generate a personalized summary using two configurations:

- **Without Ontology:** Personalization is performed using only the raw patient data (no symbolic reasoning or SWRL-based inference).
- **With Ontology:** Personalization leverages ontology reasoning, SWRL rules, and symbolic-to-natural mapping to inform LLM prompting.

On one hand, We compare the generated summaries against the original WebMD summaries (serving as a general reference) using a set of automatic evaluation metrics, and we analyze how personalization differs across configurations.

On the other hand, we use LLM-as-a-judge paradigm[145], where a secondary model evaluates the outputs of the first LLM. The idea is to assess the reasoning explanations of the LLM with metrics of coherence and completeness.

As for the SWRL rules, we defined in addition to the actual rules, some fake rules (eg: "Patient needs alien language"), in order to evaluate how the LLM interprets them, and their influence in the Personalized output.

Experiments also have been conducted regarding the hyperparameters of the retrieval aspects of the RAG layer, we focused on the issue of missing context by varying chunk sizes.

Evaluation Metrics

We adopt both text overlap and semantic evaluation metrics to assess the personalization quality. For the first evaluation, we used the following metrics [146, 129] :

- **ROUGE-1, ROUGE-L:** Measures lexical overlap (unigrams and longest common subsequences) between generated summaries and reference summaries.
- **Novelty:** Assesses the proportion of newly generated content not present in the source text, indicating the ability to rephrase or add explanatory content.
- **Coverage:** Measures how much of the essential medical content from the original article is retained in the personalized summary.
- **Relevance:** Evaluates whether the personalized summary matches the specific needs of the patient profile, based on keyword and concept matching.

As for the second evaluation, with LLM-as-a-judge [145], by defining 4 dimensions, each scored on a scale of 0 to 5 :

- Consistency : Does the reasoning explanation align with the personalization factors ?

Score	Description
0	Completely non-compliant; contradicts or ignores personalization factors.
1	Poorly compliant; violates several personalization factors.
2	Partially compliant; respects some factors.
3	Mostly compliant; minor inconsistencies with the personalization aspects.
4	Highly compliant; small deviations or minor ambiguities only.
5	Fully compliant; aligns with all relevant personalization factors.

Table 3.2 – Consistency Scoring Criteria

- Faithfulness : Is the reasoning grounded in the actual patient profile and personalization context (not hallucinated)?

Score	Description
0	Completely irrelevant; does not address the query or consent context.
1	Minimally relevant; contains some related terms but misses the main idea.
2	Partially relevant; touches on the topic but lacks essential information.
3	Moderately relevant; captures the main theme but omits key details.
4	Highly relevant; mostly aligns with the expected content, with minor omissions.
5	Perfectly relevant; fully answers the prompt or consent context precisely and coherently.

Table 3.3 – Faithfulness Scoring Criteria

- Coverage : Does the reasoning address all relevant factors rather than only a subset?

Score	Description
0	No coverage of the required content.
1	Minimal coverage; misses almost all expected points.
2	Partial coverage; includes some but not all required aspects.
3	Moderate coverage; most essential aspects are present but lacks depth.
4	High coverage; nearly all required content included with minor omissions.
5	Complete coverage; exhaustive and well-structured representation of all expected elements.

Table 3.4 – Coverage Scoring Criteria

- Clarity : Is the reasoning clear, understandable, and well-structured for human interpretation?

Score	Description
0	Completely unclear; incomprehensible or incoherent.
1	Very hard to understand; full of ambiguities or poor phrasing.
2	Understandable with effort; confusing or verbose structure.
3	Fairly clear; main ideas are understandable but not fluidly expressed.
4	Clear and well-structured; easy to follow, minor stylistic issues.
5	Exceptionally clear and readable; suitable for patient-level comprehension.

Table 3.5 – Clarity / Comprehensibility Scoring Criteria

3.4.5 Results

Figure 3.3 illustrates the comparative performance of personalized text generation with and without ontology-based reasoning. For visualization purposes, we present results related to 30 summaries for 6 patients, with 360 output generations in total. The results are reported separately for generations produced without ontologies (Figure 3.3a) and with ontology-based reasoning (Figure 3.3b).
We observe that ontology-enhanced configuration leads to significant improvements across several metrics.

Specifically, the graph in figure 3.3b shows higher values for **relevance**, **coverage**, and **novelty**, indicating that the integration of patient-specific semantic rules improves both the informativeness and the adaptability of the output.

The increase in ROUGE scores suggests that content fidelity is also preserved or enhanced. Particularly, higher novelty scores are achieved without sacrificing coverage or relevance, highlighting that the personalized content remains faithful to the original while being adapted to the patient's profile.

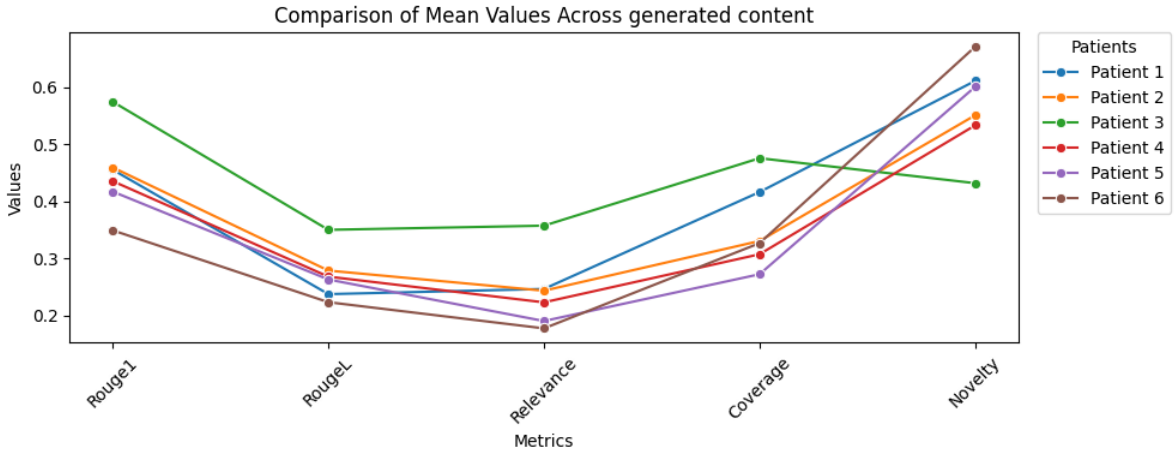
In a few cases, a slight decrease in performance is observed. This is primarily due to a mismatch between the patient's profile and the content of the source summary, where limited contextual overlap reduces the potential for meaningful personalization. Nonetheless, these results overall confirm the benefit of incorporating symbolic reasoning and structured domain knowledge in guiding large language models for patient-centered medical content generation.

In terms of the evaluation for the LLM-as-a-judge assessment, the figure 3.4 illustrates the different results for all patients across different disease categories, for the criteria (Consistency, Faithfulness, Coverage, Clarity), showing that the majority scores cluster in the upper range (3-5), which indicates that the reasoning behind the LLM generated output is of good quality.

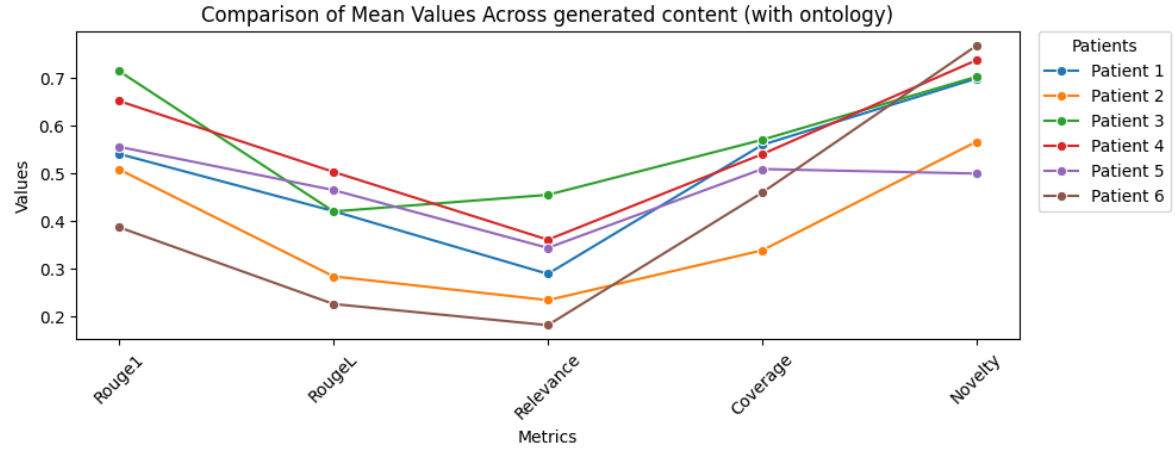
Indeed, most datasets achieve a consistency score higher than 4, which suggests that the reasoning explanations are consistent with the personalized factors. This implies that the model is generally robust, even with occasional misalignment (that are related to high divergence of the summary text and the patient's profile).

Faithfulness shows also good results but with more cases below 3, which is related to some of the LLMs hallucinations, but are very limited, and corrected by the KG rules. In terms of coverage, The results show that the LLM addresses all relevant personalization factors in its generation (with a score higher than 4). And clarity scores remain consistent (mostly 4-5), which reflects the LLM's strength in comprehensibility and explainability.

The violin plot in figure 3.4 shows also the variability of results depending on the disease category, while it suggests the LLM's dependability on the topic's complexity and it's compatibility to patient profile, it does on the other hand showcase the personalization variability



(a) Personalization without ontology



(b) Ontology-enhanced personalization

Figure 3.3 – Performance comparison of text personalization without and with ontology guidance.

for each profile.

Concerning the introduced Fake SWRL rules given in the prompt, we observed that the LLM effectively ignore them, especially when they are contradictory to the prior rules. This is due to the induced reasoning symbolic reasoning layer.

It is to be noted also that the experiments covered both models of Llama and Deepseek. But we discovered quick hallucinations with Deepseek, which made us stop the experiment. as such all the presented results are of the Llama model.



Figure 3.4 – LLM-as-a-judge on reasoning performances

3.5 Conclusion

This chapter presented the second major component of the SPIRIT framework — the semantic personalization of informed consent. Building upon the legally grounded foundation established in the previous chapter, we proposed an ontology-driven approach that integrates structured reasoning with the generative capabilities of large language models (LLMs). The central objective was to adapt consent content dynamically to individual patient profiles while ensuring transparency, explainability, and compliance with ethical and legal standards.

We began by reviewing the state of the art in LLM-based personalization systems and profile ontologies, highlighting the key limitations of existing approaches in reasoning capacity, semantic grounding, and interpretability. To overcome these challenges, we introduced a multi-layered semantic architecture composed of four layers: retrieval, semantic path extraction, reasoning, and dynamic prompt engineering. This architecture leverages the Clinical Smell and Taste Disorders Ontology (CSTDO) to represent patient profiles and consent-related concepts formally. Using SWRL-based reasoning, it infers explicit personalization needs that are then translated into structured prompts guiding the LLM's generative process.

The experimental evaluation confirmed that integrating ontological reasoning significantly enhances personalization quality, semantic coherence, and contextual relevance. Compared to baseline LLM and RAG architectures, our approach demonstrated improved results across relevance, coverage, and interpretability metrics, validating the benefit of combining symbolic reasoning with neural generation for personalized, patient-centered consent.

Having established the mechanisms for semantic personalization, the next chapter will address the third component of the SPIRIT architecture — the medical recommendation system.

Chapter 4

Care Recommendation

Sommaire

4.1 Introduction 95
4.2 Ontology-Driven Care Recommendation 96
4.2.1 Related Works 97
4.2.1.1 Medical Knowledge Graphs 97
4.2.1.2 User Profiles in Personalized Healthcare 98
4.2.1.3 Multi-Label Classification for Complex Medical Decision-Making 98
4.2.1.4 Synthesis 99
4.2.2 Care recommendation and planning system 99
4.2.2.1 Input Layer 101
4.2.2.2 Parameter Augmentation 102
4.2.2.3 Building the Knowledge Graph from Local Data 103
4.2.2.4 Multi-label classification 105
4.2.2.5 Classification model 106
4.2.2.6 Interpretation to patient 106
4.3 Experiments and Results 106
4.3.1 Knowledge Graph Construction and Analysis 107
4.3.2 Multi-Label Classification Performance 108
4.3.3 Use case: Personalized care plan 109
4.4 Conclusion 110

4.1 Introduction

Unlike medication prescriptions or clinical diagnoses, nursing care actions in home-care institutions—such as hygiene assistance, wound care, mobility support, or monitoring of vital signs—are often recorded in free text and lack any formal structure or standardized vocabulary. In daily practice, similar care activities may be described using different terms across institutions, or even by different professionals within the same service.

This absence of shared standards creates several obstacles: (1) poor interoperability between home-care institutions and information systems, (2) limited integration of nursing documentation into clinical decision-support tools, and (3) the inability to construct reliable, large-scale datasets suitable for machine-learning–based care recommendation. These challenges are further amplified in home-care settings, where patient situations are highly heterogeneous and where existing datasets are often small, incomplete, or inconsistently annotated.

As a result, care planning in home-care services frequently relies on manual interpretation of textual reports and professional experience, making it difficult to systematically reuse past cases or to ensure consistency across care teams. This limits the potential for data-driven support in proposing personalized and adaptive care plans, particularly for elderly or multi-pathological patients receiving long-term home care.

This chapter addresses these limitations by proposing an ontology-driven, multi-label care recommendation approach that formalizes nursing standards and links them to individual patient profiles. By structuring home-care activities—such as personal care, mobility assistance, therapeutic monitoring, or preventive interventions—within a shared semantic framework, the proposed approach provides the foundation required for interoperability, explainable recommendation, and predictive modeling.

With this chapter, we present the final layer of the SPIRIT architecture, extending the previous contributions on legal compliance and personalized consent by ensuring that consent is not only compliant and understandable, but also operationally connected to individualized care recommendations within real-world home-care practices.

4.2 Ontology-Driven Care Recommendation

This section focuses on the third core contribution of the SPIRIT framework: a semantic and data-driven approach for recommending nursing care actions in home-care institutions. The motivation for this component arises from a major gap in current practice: nursing care actions are neither standardized nor encoded in a structured manner. As a result, institutions produce heterogeneous, non-interoperable documentation that hinders reasoning, analytics, and automatic recommendation. In that matter, we introduce in section 4.2.1 an overview of related works on medical Knowledge graphs, user profiling and multi-label classification for decision making.

Along With this issue, the limited availability of clean, labeled datasets makes classical machine learning approaches impractical. To address these challenges, we introduce in section 4.2.2 an ontology-based standardization of nursing actions, linked with patient-specific data, and a multi-label classification model enhanced through parameter augmentation. This component evaluates how semantic structuring and predictive modeling can jointly support reliable, personalized care recommendations personalized to home-care contexts. Finally, experimental results are showcased in section 4.3, focusing on the whole performance of the recommendation system.

4.2.1 Related Works

Personalized care planning is increasingly seen as a key objective in modern healthcare systems. Yet in practice, home-care institutions often struggle to provide decision support due to heterogeneous data formats, non-standardized nursing actions, and incomplete patient records. Artificial Intelligence (AI) has shown promise in addressing these challenges by extracting patterns from electronic medical records, sensor data, and clinical narratives [147]. However, both the *heterogeneity* and *incompleteness* of these data sources remain persistent obstacles to robust and explainable care recommendation.

To situate our contribution within the broader research landscape, this section reviews three interconnected areas: (1) Medical Knowledge Graphs (MKGs) and their role in structuring clinical data and enabling semantic inference; (2) User profiles and profile ontologies as foundations for personalized healthcare; and (3) Multi-label classification (MLC) approaches for modeling complex, multi-pathology medical decisions.

4.2.1.1 Medical Knowledge Graphs

Medical Knowledge Graphs (MKGs) provide a powerful representation framework for integrating heterogeneous clinical data, capturing semantic relations between diseases, symptoms, interventions, and patient states. They have been used to support diagnosis, drug repurposing, risk prediction, and treatment personalization.

Henaien et al. [147] illustrate this trend through a system that merges patient sensor data with ontologies such as ICNP and FOAF into a unified knowledge model. Reasoning is implemented through SWRL rules, whose outputs guide a decision-tree–based predictive engine. Although effective, the approach remains limited by its reliance on hand-crafted rules and its sensitivity to noisy or missing information [148].

A comprehensive survey by Wu et al. [148] identifies three main MKG reasoning paradigms: (i) semantic pattern–based reasoning, (ii) symbolic reasoning directly over graph structures, and (iii) knowledge representation learning (KRL) models. These paradigms reveal the trade-offs between interpretability and scalability. Symbolic reasoning provides transparency but struggles with incompleteness; embedding-based models scale better but lose semantic explainability.

Recent studies show that embedding-enhanced methods [149, 150] can improve robustness and predictive accuracy, especially in complex clinical scenarios. However, their reliance on vector representations makes it difficult to provide traceable justifications — a key need in healthcare recommendation.

4.2.1.2 User Profiles in Personalized Healthcare

Effective recommendation requires a precise understanding of the patient, operationalized through the concept of a user profile. In healthcare, profiles include demographic attributes, medical history, treatment preferences, social determinants of health, and contextual environmental factors [151].

The construction of such profiles is hindered by heterogeneous data schemas and the absence of shared semantics. Ontologies address this issue by providing structured, interoperable, machine-interpretable representations of personal and clinical attributes [55]. We follow on our state of the art on user profile ontologies introduced in section 3.2.2, that covered different ontologies : Foaf [135], HL7 [136], Snomed CT [7], E-PPO [137], CSTDO [138], HSPO [139], CRENO [140].

In our work, we rely primarily on the Clinical Smell and Taste Disorders Ontology, not for its biomedical specificity, but because it provides structured demographic and personal descriptors closely aligned with those found in our real-world dataset. Which also have been used in the personalization component of SPIRIT architecture.

4.2.1.3 Multi-Label Classification for Complex Medical Decision-Making

Healthcare decisions often involve simultaneous conditions and multiple recommended actions. Multi-label classification (MLC) methods are therefore essential for predicting care interventions that are interdependent and context-specific. Applications appear in disease prediction, symptom grouping, and treatment selection [152, 153].

MLC approaches fall into three families [154]:

- **Problem Transformation**: e.g., Binary Relevance, Classifier Chains.
- **Algorithm Adaptation**: e.g., ML-kNN, ML-SVM.
- **Ensemble Methods**: enhancing robustness through model combination.

Recent models integrate MKGs and MLC. MulGral [155], for example, uses a graph-based framework to jointly model patients, symptoms, and diseases, outperforming traditional MLC models. Likewise, KGAT [150] demonstrates how attention mechanisms over graph embeddings can improve relational interpretability in multi-label prediction.

More recently, LLM-based recommenders [156, 157] have emerged. While LLMs excel at leveraging linguistic patterns, they struggle with factual grounding and may hallucinate recommendations. This has motivated hybrid approaches combining LLMs with knowledge graphs [158], leveraging the explicit semantics of graphs to reduce hallucinations and increase explainability.

4.2.1.4 Synthesis

The literature highlights the strengths and limitations of the three main domains relevant to our work:

- **MKG-based reasoning** offers semantic structure and interpretability, but symbolic methods are brittle and embedding methods lose transparency.
- **Profile ontologies** provide a solid foundation for interoperable and semantically rich patient representations, yet are rarely integrated into predictive pipelines for home-care contexts.
- **Multi-label clinical decision models** can predict complex care actions, but require structured datasets and often lack semantic grounding, especially when nursing actions lack standardization.

Three critical gaps emerge:

1. **Lack of standardized nursing action vocabularies**, making it difficult to build datasets suitable for multi-label prediction.
2. **Limited integration of patient semantic profiles into predictive models**, restricting personalization capabilities.
3. **Insufficient explainability in embedding- and LLM-based methods**, which is problematic for clinical decision support.

These gaps motivate our approach, which:

- formalizes nursing care actions through an ontology-based care standard,
- models patient profiles using the CSTDO ontology to ensure structured, interoperable personalization, and
- integrates these semantic representations into a multi-label classification model enhanced with parameter augmentation.

This combination provides the foundations for scalable, explainable, and personalized care recommendations tailored to real home-care institutional workflows.

4.2.2 Care recommendation and planning system

Clinical datasets used in local care institutions typically lack standard terminologies, abstraction levels, and semantic interoperability. The dataset available in our study originates from a

Approach	Study	Knowledge Representation	Transformer Use	User Profiling / Personalization	Evaluation / Results	Limits
Ontology-based ML for medical decision	Henaien et al. (2020)	Semantic rules using ICNP, FOAF, etc.	No	Yes (user profile DB)	Decision tree logic	Rule generation intensive
	Shi et al. (2017)	MeSH, UMLS ontologies in a semantic KG	No	Yes (health services recommendation)	Binary classification (custom)	High processing cost
KG + ML/DL for multi-label prediction	Pham et al. (2022)	Ontology-based medical KG	Optional	Yes (disease prediction per patient)	Multi-label classification	Complex KG structure
	Chai et al. (2020)	KG embeddings from thyroid data	Yes (DL)	No	Deep learning outperforming ML baselines	Domain-specific to thyroid
KG-enhanced LLMs	Zhao et al. (2024)	KG with LLM-based extractor (E-E-I recall)	Yes	Yes (profile-based KG traversal)	AUC, Click, GMV	Long user history performance issues
	Shu et al. (2024)	KG + LLM for link prediction	Yes	No	Hits@1/3/10	Costly inference, less reasoning depth
LLM-enhanced profiling	Zhou et al. (2024)	No KG (LLM vector space)	Yes	Yes (LFM: explicit profile)	RMSE, MAE, Choice Error	Fails with long history, poor generalization
Synergistic LLM–KG	Yang et al. (2023)	Implicit KG + interaction graph + LLM planner	Yes	Yes (3-stage personalization)	HR@10, NDCG@10	Computational cost; static profiles

Table 4.1 – Summary of state-of-the-art investigations on data representation, ontology usage, transformers, classification types, and results achieved in medical decision systems

non-standardized nursing information system and includes patient characteristics, diagnoses, and care actions, mostly expressed in free text. This introduces two limitations: (1) learning models cannot be trained on reliable semantic representations, and (2) recommendations cannot be generalized beyond this dataset because no standardized care knowledge exists.

We address these limitations by building a complete pipeline that:

- standardizes local terminology using international nursing ontologies,
- constructs a knowledge graph (KG) connecting patients, diagnoses, and care actions,
- embeds the KG in vector space for learning,
- predicts care recommendations through multi-label classification,
- generates human-interpretable explanations of the recommendations.

Figure 4.1 illustrates the architecture of the proposed approach. It contains four layers :

- Input : Contains the different data sources used, the local patient dataset, ICNP taxonomy and Snomed ontology.
- Parameter augmentation : Serves at creating a knowledge graph from the mappings between the different sources, through an LLM, which adds more parameters to the whole dataset.
- Multi-label classification : Aims at creating a classification model that is trained on the augmented KG of the previous layer, through a TransE embeddings.
- Output : The resulted classification model is used to predict care actions for a given patient. The prediction (recommendation) is then given to the patient in natural language via an LLM.

4.2.2.1 Input Layer

In the architecture's input process, we work with three main data sources : the local patient dataset of the health institution (Léonie Chaptal Foundation), ICNP, and Snomed CT.

- **The International Classification for Nursing Practice (ICNP)** [159]: The ICNP provides a formalized nursing vocabulary organized around a seven-axis model (*Focus, Judgment, Means, Action, Time, Location, Client*). It defines relationships such as `hasFocus`, `hasJudgment`, `performedBy`, `performedAt`, and `hasMeans`, allowing fine-grained representation of clinical activities. Later releases expanded the terminology to around 4,000 concepts and approximately 30 semantic relations. Importantly, ICNP supports direct and partial mappings to SNOMED CT, covering up to 83% of its concepts.

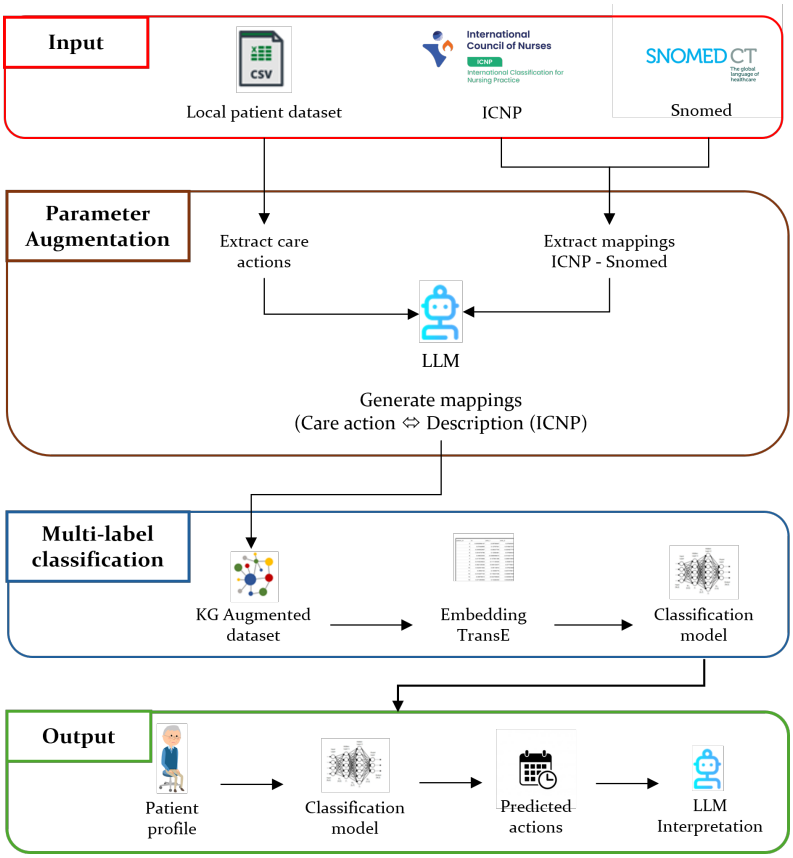


Figure 4.1 – Care Recommendation and planning system.

- **SNOMED CT** [7]: SNOMED CT contains over 350,000 active concepts and 1.3 million semantic relationships, including core relations such as `is_a`, `part_of`, `associated_with`, and `has_causative_agent`. It provides the foundational clinical backbone for interoperability, allowing integration of nursing, diagnostic, and procedural data under a shared semantic layer.

4.2.2.2 Parameter Augmentation

This layer aims to transform heterogeneous, text-based nursing data into a structured, interoperable knowledge graph capable of supporting personalized care recommendations. Our dataset consists of real-world patient profiles, diagnoses, and recorded care procedures. However, these care actions were initially documented using non-standardized, institution-specific terminologies, which limited semantic interoperability, reasoning, and reuse across systems. To address this limitation, we applied a semantic normalization process that aligns local data with internationally recognized nursing ontologies.

ICNP as the core nursing taxonomy. It provides the conceptual vocabulary for nursing actions, judgments, and care procedures. We reused concepts and relations from the following

axes:

- Action / Care Actions: e.g., Bathing patient, Monitoring vital signs, Pressure injury prevention.
- Focus / Condition: e.g., Mobility, Skin integrity.
- Judgment / Status: e.g., Decreased mobility, Pain present.
- Means / Resources: e.g., Assistive devices, Medication.
- Relations reused: `hasFocus`, `hasMeans`, `hasJudgment`, `performedAt`, `performedBy`.

SNOMED CT for clinical diagnosis and hierarchical reasoning. It supplies medical and diagnostic background:

- Clinical findings and diseases: Hypertension, Diabetes, Pressure ulcer.
- Body structure and clinical context.
- Relations reused: $\texttt{is}_a$, `partOf`, `associatedWith`, `causativeAgent`.

ICNP and SNOMED are therefore complementary: ICNP models the care action domain, while SNOMED models conditions and diagnosis.

New Concepts and Relations Introduced Because neither ICNP and Snomed CT directly models the recommendation setting, we add a minimal schema layer:

- New classes: `Patient`, `Diagnosis`, `CareAction`, `ProfessionalRole`, `CareSetting`
- New object properties: `hasDiagnosis` (Patient → Diagnosis) `requiresCareAction` (Diagnosis → CareAction) `hasProfileAttribute` (Patient → Attribute)

This layer explicitly connects patient profiles with ICNP/SNOMED elements.

4.2.2.3 Building the Knowledge Graph from Local Data

To create the augmented KG, we follow several steps :

- **Step 1: Text preprocessing :**
  Local nursing procedures and patient profiles are cleaned and structured: extraction of diagnoses and care actions, sentence segmentation.

- **Step 2: Ontology Mapping using LLMs :**
  Traditional string similarity is insufficient due to linguistic variability. We therefore perform LLM-assisted semantic alignment:
  - The model receives: the local care phrase, contextual information, a list of candidate ICNP/SNOMED concepts, definitions and hierarchical parents. The context and the prompt given to the LLM is presented below.
  - It outputs RDF/Turtle triplets expressing alignment and confidence.

  Prompt given to the LLM:

```
Role: You are a semantic similarity mapping agent.
Goal: Map each nursing term from a local dataset to the adequate ICNP
    ontology concept
based on their meaning, context, and definitions.

Local nursing term: "Bed bath"
Context: "Daily hygiene assistance provided to bedridden patients."
Candidate ICNP concepts:
  - icnp:BathingPatient ("Cleansing the body of a patient")
  - icnp:ComfortCare ("Providing comfort to the patient")
  - icnp:SkinAssessment ("Assessing the condition of the skin")

Output:
Mappings in JSON-LD or Turtle format with confidence scores and reasoning
     justification.
```

  The LLM produces a structured output in TTL format that encodes explicit mappings between the local and ontology concepts. Each mapping is annotated with metadata capturing:
  - A **confidence score** (ranging from 0 to 1), quantifying the semantic alignment strength between the two entities.
  - A **rationale** summarizing the semantic justification for the mapping, i.e., why the model believes the two terms are equivalent.
  - Optionally, **alternative candidates** with lower confidence, which can be used for human validation or ontology enrichment.

  Example of a generated mapping:

```
ex:BedBath a icnp:BathingPatient ;
    rdfs:label "Bed bath" ;
    ex:confidenceScore "0.92" ;
    ex:alignmentRationale "Action involves cleaning the patient using
    water while bedridden;
    equivalent to 'BathingPatient' in ICNP.";
    ex:alternativeCandidate icnp:ComfortCare ;
    ex:confidenceAlt "0.61" .
```

  The alignment process thus produces a comprehensive, machine-readable mapping file in RDF/Turtle format, which can be directly ingested into the knowledge graph. These

alignments are not static, in which new local terms added to the dataset automatically trigger re-alignment by the LLM.

The integrated graph consequently unifies local clinical data and standardized nursing terminologies under a shared semantic model. ICNP serves as the primary ontology for care actions and nursing judgments, while SNOMED CT provides hierarchical and causal relations (`is_a`, `associated_with`, `has_causative_agent`), enriching inferential capacity. This structure enables reasoning across related care concepts, supports automated recommendation generation, and ensures interoperability with other healthcare information systems.

Each mapping contains: matched ontology entity, confidence score, semantic explanation, alternative candidates. Mappings are generated locally using Llama-3.1-8B.

- **Step 3: Creation of the Final Knowledge Graph :**
  The aligned terminology is merged with ICNP and SNOMED fragments. The resulting KG contains: 880 nodes, 1,733 edges.
  Type of entities: - patients, diagnoses, care actions, - ICNP concepts, - SNOMED concepts, - settings, roles.

  Relations:
  - hierarchical and semantic relations from the ontologies, - new relations linking patients diagnosis care actions.

  The KG is represented in RDF/OWL and stored as triples.

4.2.2.4 Multi-label classification

In this layer, we focus on creating a classification model through an embedding the augmented Knowledge Graph. To prepare this knowledge graph for learning tasks, we embed both entities and relations into a continuous vector space using knowledge graph embedding methods. We adopt the TransE model, which represents each triplet (s, r, t) (source entity, relation, target entity) as vectors $(\vec{v}_s, \vec{v}_r, \vec{v}_t)$ and optimizes the following relation:

$$\vec{v}_s + \vec{v}_r \approx \vec{v}_t$$

The objective is to minimize the distance:

$$\|\vec{v}_s + \vec{v}_r - \vec{v}_t\|$$

for valid triplets and maximize it for corrupted ones. These embeddings capture the structural and semantic properties of each entity based on its local and global neighborhood in the graph.

As a result, patient profiles become vectorized representations that encode clinical context, inferred relations, and semantic proximity between diagnoses and care actions.

4.2.2.5 Classification model

With enriched vector representations, we train a multi-label classification model to predict relevant nursing procedures for each patient. We use a Multilayer Perceptron (MLP) architecture that takes as input the patient embedding (of size 202) and outputs a probability vector over 36 possible care actions:

$$\{p_1, p_2, \ldots, p_{36}\}$$

The network includes:

- A first hidden layer with 128 neurons,
- Intermediate hidden layers with 64 neurons,
- An output layer of 36 neurons with sigmoid activation.

Dropout (rate = 0.3) is applied after each hidden layer to prevent overfitting. The use of sigmoid activation supports multi-label prediction, enabling the simultaneous recommendation of multiple care actions per patient.

4.2.2.6 Interpretation to patient

To enhance interpretability and trust, the predicted care actions are passed to an LLM (with the same configuration previously employed : Llama 3.1 8b), which generates a textual explanation contextualized by the patient's profile. While LLMs are known to hallucinate, grounding their outputs in the knowledge graph provides factual constraints and domain-specific structure.

By integrating ontological facts and associations into the LLM prompt (e.g., via retrieval-augmented generation), the model is able to produce logical, patient-specific justification for each suggested procedure. This step enables patients and clinicians to better understand the rationale behind the system's suggestions, facilitating shared decision-making and transparency.

In short, instead of delivering the predicted action as a concept, the LLM is asked to clarify it and explain it, based on it's ontological description. To prevent the LLM from hallucination, the ontological layer of ICNP is added, with the concept's description, which limits it's generation to the scope of the ontology.

4.3 Experiments and Results

In this section, we evaluate the proposed framework across three key layers: (1) the construction of the knowledge graph from real-world patient data, (2) the performance of the

multi-label classification model, and (3) the integration of LLMs for personalized and explainable recommendations. Our main objectives are as follows:

- **G1** Build a structured medical Knowledge Graph from patient data and standard ontologies.
- **G2** Recommend personalized care plans based on knowledge-aware embeddings.
- **G3** Develop and evaluate a multi-label classifier for treatment prediction.
- **G4** Assess the contribution of LLMs in generating contextual, patient-specific explanations.

4.3.1 Knowledge Graph Construction and Analysis

The knowledge graph was constructed with respect to a dataset of real patient records, including age, gender, diagnosed pathologies, and nursing treatments prescribed. According to the approach, such data points were mapped to standard SNOMED CT and ICNP terminologies. The resulting graph encodes semantic and clinical relations between three main entity types: patients (P), medical conditions (M), and care procedures (S).

In the mapping generation, we use Llama 3.1 with 8b parameters, locally imported and executed in Colab environment (with GPU), and the temperature used was fixed in 0.7.

Metric	**Value**
Nb of Nodes	880
Nb of Edges	1733
Nb of Detected Cycles	0
Nb of Weakly Connected Components	1
Nb of Leaves (Terminal Concepts)	6

Table 4.2 – Knowledge Graph Evaluation

As shown in Table 4.2, the graph contains 880 nodes and 1733 edges. Although relatively small due to the limited dataset, its topology is meaningful. The average connectivity degree (of the order of 2 relations per node) indicates a good connectivity, which is needed in knowledge graph embedding.

The graph is acyclic, with no cycles, confirming its hierarchical, acyclic structure. This supports the ontology-based design, which excludes semantic loops and ensures consistent information flow. Secondly, since there is just one weakly connected component, it ensures all entities can be reached, enhancing consistency of training signals in embedding.

4.3.2 Multi-Label Classification Performance

Using TransE, we embedded all patient profiles into a continuous vector space. These embeddings served as input to a Multi-Layer Perceptron (MLP) for predicting appropriate care actions.
Figure 4.2 illustrates the evolution of the F1 Micro, AUC Micro, and Hamming Loss across the 10-fold experiments. The global average metrics are summarized in Table 4.3.

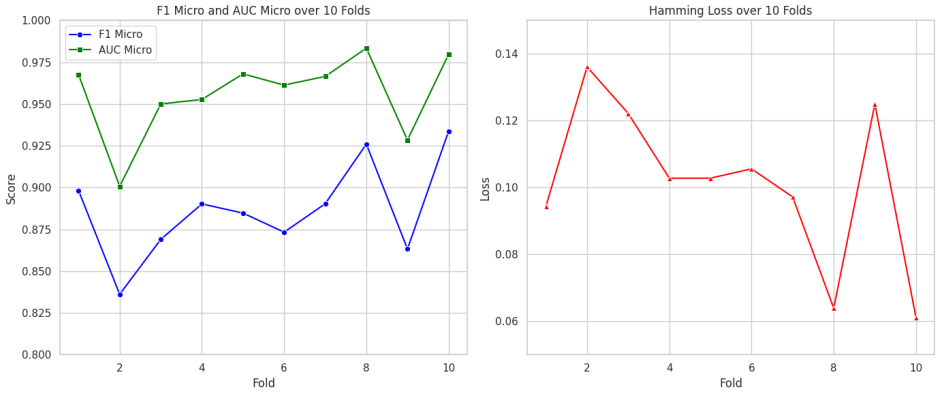


Figure 4.2 – Performance metrics across 10-fold cross-validation: (Left) F1 Micro and AUC Micro; (Right) Hamming Loss

Metric	Value
F1 micro	0.8960
AUC micro	0.9662
F1 macro	0.8960
AUC macro	0.7981
Hamming Loss	0.0936

Table 4.3 – Multi-Label Classification Results

The classifier's micro-average F1 score (0.8960) and AUC (0.9662) reflect extremely high overall recall and precision on all the treatments predicted. This reflects the model's ability to predict well when it is classifying patients as a group.

The macro-average F1 and AUC scores confirm that the model also handles label diversity effectively. With AUC macro = 0.7981, the system distinguishes even rare treatment cases — a common challenge in medical datasets with class imbalance. The F1 macro value remains high at 0.8960, indicating that the model is not biased toward frequent classes. Hamming Loss of 0.0936 implies a low per-label error rate: fewer than 10% of treatment predictions are incorrect per patient.

Despite the good results, the slightly lower macro AUC highlights some difficulty in distinguishing rare cases. This could be mitigated in future iterations by balancing the training set or applying label-aware loss functions.

4.3.3 Use case: Personalized care plan

The output of the supervised classifier and the knowledge graph is a set of recommended procedures aligned with the patient's profile (e.g., age, pathology, autonomy level, and care setting). While the classifier provides structured predictions, communicating these recommendations to clinicians and patients in an interpretable manner is still non-trivial. To address this challenge, we integrate a Large Language Model (LLM) to generate personalized natural-language explanations for the predicted care plan.

The LLM operates over two sources of structured knowledge: The predicted procedures obtained from the multi-label classifier, expressed as ICNP/SNOMED concepts. And the factual constraints stored in the knowledge graph, describing clinical motivation, applicability conditions, and semantic relations between concepts.

The LLM receives as input a structured prompt containing:

- the patient's structured profile;
- the predicted procedures;
- the contextual triples retrieved from the KG through semantic querying.

This allows the LLM to generate a textual explanation that is clinically grounded rather than purely generative. The KG plays two major roles in this process: It constrains the generation to concepts present in the formal ontology (procedure types, contraindications, clinical associations). And it provides a factual scaffold that reduces hallucinations by forcing the LLM to justify recommendations based on true, source-validated statements.

This hybrid process transforms the system from merely predicting a procedure toward explaining why that recommendation is appropriate for this specific patient.

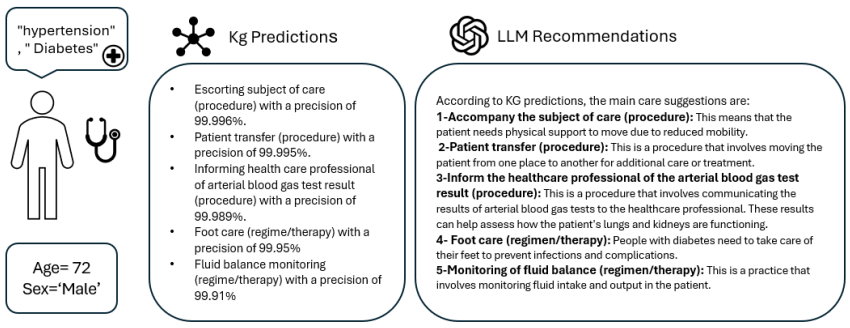


Figure 4.3 – Use case: Personalized care plan

Figure 4.3 illustrates this process. Given a patient diagnosed with diabetes and reduced autonomy, the classifier recommends procedures such as Nutrition monitoring, Medication

supervision, and Fall-risk prevention. The knowledge graph augments these predictions with semantic justification (e.g., diabetes → nutrition control; advanced age → fall prevention strategies), retrieved through ontology-based reasoning.

The LLM then converts these facts into contextualized explanations, explicitly linking each procedure to the patient's condition. For example: "Nutrition monitoring is recommended because diabetes requires regular control of glycemic intake."

Through this chain, each recommendation becomes traceable through three layers:

- A data-driven prediction with the LLM mappings and classification,
- A semantic justification through the knowledge graph,
- A human-readable explanation with the LLM.

This hybrid architecture demonstrates a transition from prediction to interpretable and clinically motivated decision support, bridging structured medical ontologies with natural-language communication. It improves transparency, supports clinical validation, and increases trust and adoption in care-planning environments.

4.4 Conclusion

In this chapter, we addressed the limitations inherent in local, non-standardized nursing datasets and proposed a hybrid semantic–machine learning approach to improve care recommendation in home-care institutions. By restructuring raw clinical data through ICNP and SNOMED CT mappings, and enriching it with ontology-driven contextual knowledge, we created a more coherent and interoperable representation of patient conditions and care actions. This semantic augmentation proved essential for strengthening the performance of multi-label classification models, compensating for data sparsity, and ensuring that recommendations align with established nursing standards.

The results demonstrate the value of integrating structured medical knowledge with predictive modeling: the system not only generates more accurate and comprehensive recommendations but also improves interpretability by grounding its outputs in well-defined clinical concepts. Furthermore, the use of hierarchical relations and domain semantics enables a recommendation process that is both consistent with professional nursing practices and adaptable to different institutional settings.

Overall, this contribution provides a foundation for standard-compliant care planning in real-world environments. It sets the stage for a foundation of a proof of concept (Appendix A). There, we integrate personalized consent generation and ontology-based care recommendation into a unified workflow, demonstrating how these components operate within an actual home-care institution's processes.

Chapter 5

Conclusions et perspectives

Introduction

This final chapter provides a synthesis of the work presented in this thesis. we present first in section 5.1, a summary as well as some concluding remarks regarding the different contributions of the thesis. Then, we discuss some limitations of this work in section 5.2. Finally, future directions are described in section 5.3.

5.1 Summary

This thesis made novel contributions to semantic personalization via knowledge graphs and LLM integration applied in the context of informed consent. At least four main contributions can be identified in our work : (i) introducing a new modular approach for informed consent (Chapter 2), (ii) Providing an ontology-driven machine learning approach for legal compliance (chapter 2), (iii) Proposing a semantic reasoning mechanism integrating ontological knowledge with LLM generation for personalized consent (chapter 3), (iv) Designing a multi-label classification model for nursing recommendation, extending consent personalization to actionable care guidance (chapter 4), and implementing a proof of concept web application demonstrating the integration and practical usability of the overall approach (chapter 4). These contributions have answered the research questions defined in the beginning of this work.

The main objective of this thesis is to design and implement a semantically grounded system that ensures compliance, personalization, and care recommendation in the generation and management of informed consent documents within a healthcare institution. To achieve this goal, we propose SPIRIT (Semantic Personalized and Compliant System for Informed consent and Care recommendation), a global architecture that integrates knowledge graphs, semantic reasoning, and LLMs to support three main pillars of the consent process: legal compliance, consent personalization, and care recommendation. The approach leverages ontological structures to ensure interpretability and traceability of information, while machine learning and natural language generation modules provide adaptability and scalability. The SPIRIT system thus operationalizes a continuum from compliant consent documentation to personalized patient communication and informed nursing recommendation, forming a unified semantic and intelligent ecosystem for medical decision support.

The first contribution addresses the problem of ensuring the automatic compliance of consent documents with evolving legal and ethical standards. Despite the increasing digitization of medical records, healthcare institutions still face challenges in verifying that consent documents meet the formal and content requirements mandated by regulations such as the GDPR and the Oviedo Convention. A synthesis of related works reveals that existing NLP approaches for legal document analysis primarily rely on text embeddings or rule-based systems, often lacking the semantic interpretability required for legal reasoning and domain traceability.

To overcome these limitations, we propose an Ontology-Driven Machine Learning Annotation approach, which integrates a legal ontology with machine learning models for semantic document annotation. The ontology encodes legal norms, obligations, and document types, serving as a semantic layer that guides the annotation process and enhances model interpretability. Experiments were conducted using transformer-based embeddings with a hyperparameter optimization, and the influence of ontological constraints was evaluated on the classification accuracy of document sections. The results demonstrate that incorporating semantic relations from the ontology improves similarity scores measurements, providing an interpretable and compliant annotation mechanism. This component establishes the legal foundation of SPIRIT, ensuring that every subsequent process operates within a legally verified and semantically coherent framework.

The second contribution focuses on the personalization of informed consent documents, addressing the need to adapt consent content to individual patient profiles. Traditional consent forms often remain generic and poorly understood by patients, especially those with low literacy or diverse cultural and linguistic backgrounds. A review of existing e-consent systems reveals that while they improve accessibility, they rarely achieve true personalization grounded in semantic and contextual knowledge. Moreover, LLMs have shown potential for content adaptation but lack structured integration with domain ontologies and reasoning mechanisms.

In response, we developed a Semantic Prompting for Consent Personalization approach, a hybrid framework integrating LLMs with ontology-based reasoning and profile knowledge graphs. Patient profiles are modeled semantically, capturing demographic, clinical, and cognitive characteristics, which are then used to condition generative models through structured prompts. The system dynamically adjusts the tone, complexity, and relevance of the content according to the inferred patient profile. Experimental evaluations were conducted to assess personalization quality using metrics such as ROUGE, Relevance, and Readability, showing significant improvements over baseline LLM generation without semantic conditioning. The approach thus contributes to bridging the gap between compliance and comprehension, ensuring that patients receive consent information that is both legally sound and personally meaningful.

The third contribution extends the SPIRIT system to clinical decision support, focusing on personalized nursing care recommendation. Once compliance and personalization are ensured, the system aims to translate the consented information into actionable guidance for healthcare professionals. Related research in nursing decision support highlights the use of

ontologies and machine learning for care planning, yet existing models often lack integration with ontology standards or fail to provide semantically transparent reasoning.

To address this, we propose an Ontology-Based Multi-label Classification model for Nursing Recommendation, which combines patient data, clinical ontologies, and care pathway knowledge to generate individualized care plans. The system learns from annotated patient datasets, linking symptoms, pathologies, and care needs to appropriate nursing interventions. The proposed solution not only aims to respect ontology standards in medical field, but also to increase the parameter spectrum of local data, which helps in the classification model. The approach has been evaluated on one hand, regarding the created graph by detecting the number of nodes, concepts and cycle in the graph. And in the other hand, the classification model is evaluated with metrics of precision and recall. The approach thus reflects the third pillar of the SPIRIT system, providing assistance to health professionals in their care recommendation.

A web application prototype was developed to integrate all contributions of SPIRIT, enabling healthcare professionals to visualize patient profiles, personalized consent, and corresponding care recommendations in a unified interface. Experimental validation demonstrates that ontology-enhanced models outperform purely data-driven baselines in both accuracy and explainability. The deployment of the prototype in collaboration with healthcare professionals at the Fondation Léonie Chaptal paves the way for field evaluation and real-world assessment of the framework's usability and impact.

5.2 Limitations

We should note that this thesis has some limitations in the methodologies and the capabilities of the adopted technologies. These limitations are discussed below.

The first limitation that we can note, concerns the scope of datasets used throughout the different components of the SPIRIT approach. Indeed, the synthetic data allowed for experimental validation of the contributions, yet they may not fully reflect the diversity and complexity of healthcare environment. Particularly, personalization and recommendation might benefit from larger and heterogeneous datasets.

A second limitation relates to the integration of LLMs with ontological reasoning. Despite the effectiveness of the proposed approach in personalization, there are more parameters we can explore, especially LLMs parameters and prompt engineering. There's also the question about LLM's size, particularly for our use case, as it needs to be integrated locally, which means the need for experimenting on small LLMs. That being said, we are also limited by the metrics in terms of personalization of content generated by LLMs. The literature does introduce some of them – that we used in our work – but more unified metrics can help validate our contribution.

Lastly, in regard to the tool implemented, the web application serves at the moment as

a proof of concept. Although it demonstrates the feasibility of the SPIRIT approach, it has not yet undergone extensive testing with healthcare professionals and patients. Overall, these limitations not only reflect the complexity of this research, but they also open avenues for further development.

5.3 Perspectives

Building upon the identified limitations, several research perspectives emerge to extend the proposed work.

A first direction involves conducting larger-scale experiments on real-world data from diverse clinical settings. Such experiments would make it possible to confront the SPIRIT system with the complexity of field conditions while performing systematic comparisons with state-of-the-art approaches. These evaluations would help assess the robustness, generalization capacity, and overall impact of the system, particularly regarding personalization quality, legal compliance, and clinical usability.

Another important perspective concerns evaluating the system directly in practice, by deploying it in healthcare facilities and collecting feedback from both patients and professionals. This feedback would help measure how personalization improves comprehension, trust, and adherence, while also examining how it might reduce the administrative complexities for healthcare professionals.

From a technical standpoint, several aspects deserve exploration. A natural step would be the integration of multilingual capabilities, ensuring that the system supports equitable access to medical information in multilingual contexts. The inclusion of multimodal components—such as visual or biomedical data—also opens promising opportunities, particularly for supporting patients with neurodegenerative diseases, where communication barriers are significant.

Improving explainability and transparency remains a challenge as well. Future works of the system should explore hybrid neuro-symbolic architectures, capable of making reasoning steps explicit and ensuring the traceability of LLMs personalization. This direction is key to strengthening both clinician trust and regulatory accountability.

Finally, further work should address interoperability and infrastructure. Integrating SPIRIT with existing health information systems, through standards such as HL7 FHIR, SNOMED CT, or LOINC, would be crucial for ensuring sustainability and real-world adoption. At the same time, careful consideration must be given to data hosting, storage, and governance, particularly regarding security, sovereignty, and compliance with GDPR. These reflections will be essential for deploying the system safely and responsibly within hospitals and medico-social institutions.

Appendix A

Proof of Concept

This appendix presents the proof-of-concept system that uses the SPIRIT approach within the workflow of a home-care institution. The goal of this prototype is two-fold: (1) to demonstrate how personalized consent generation and care recommendation can be integrated into real institutional processes, and (2) to validate the practical applicability of our semantic and predictive models in a usable interface.

The proof of concept models the life cycle of a patient in a home-care setting—from initial consent generation to personalized care planning. It integrates ontology-driven personalization, semantic reasoning, multi-label predictive models, and institutional nursing standards into a coherent web application. Through this implementation, we show how SPIRIT supports healthcare professionals in delivering compliant, personalized consent documents and generating patient-adapted nursing recommendations within an operational environment.

In this appendix, we present our proof of concept as a web-based application. First we showcase the system overview and architecture. And Then we present the different modules and implementation of the tool.

A.0.1 System Overview and Architecture

The prototype developed within this research serves as a proof of concept for the proposed informed consent management system. Its objective is to put in into action the proposed approach from knowledge representation to personalized consent generation and care planning assistance.

The system was implemented as a **web application** developed in *React* for the front-end interface and *Flask (Python)* for the back-end services. A *MongoDB* database was chosen for data persistence due to its flexibility in handling heterogeneous patient and document data structures. The platform provides a unified interface for healthcare professionals to manage patient profiles, generate and review consent documents, and access personalized recommendations.

The core architecture integrates three functional modules:

1. A **consent management module** for uploading, structuring, and storing consent documents;

2. A **personalization module** for adapting document content based on patient profiles;

3. A **recommendation module** for nursing care planning and scheduling.

This design facilitates scalability, interoperability with clinical systems, and future integration with hospital information infrastructures. Figure A.1 shows the dashboard of the application.

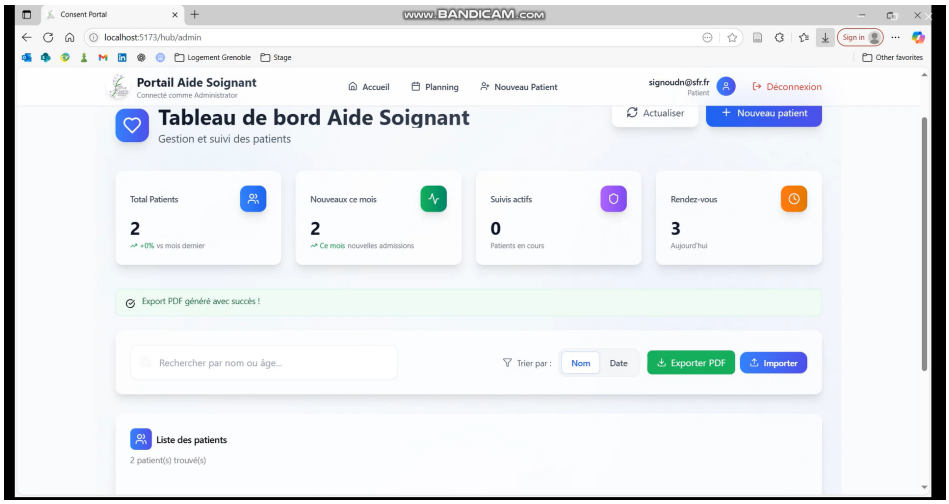


Figure A.1 – Dashboard of the application

A.0.2 Personalization of Consent Documents

A central contribution of the prototype is the integration of **semantic personalization** mechanisms, leveraging LLMs and ontology-based reasoning. An API built with *Flask* enables communication between the front-end and the LLM engine (*LLaMA 3*) for content generation and adaptation.

The personalization process operates as follows:

1. Consent documents in PDF format are uploaded and automatically segmented into semantic sections;

2. Each section is analyzed to identify key normative elements (obligations, permissions, prohibitions);

3. The LLM reformulates and summarizes the content according to the patient's profile (e.g., language, literacy, pathology, and comprehension level);

4. The resulting personalized text is then displayed on the portal in a patient-friendly format.

This architecture allows the transformation of standard legal texts into personalized, understandable, and compliant consent documents. Although hosting the LLM model introduced interoperability challenges (notably CORS issues between React and Flask endpoints), a hybrid approach—combining a local Flask server for data preprocessing and a remote server for LLM inference—was adopted successfully. Figure A.2 shows the personalization module.

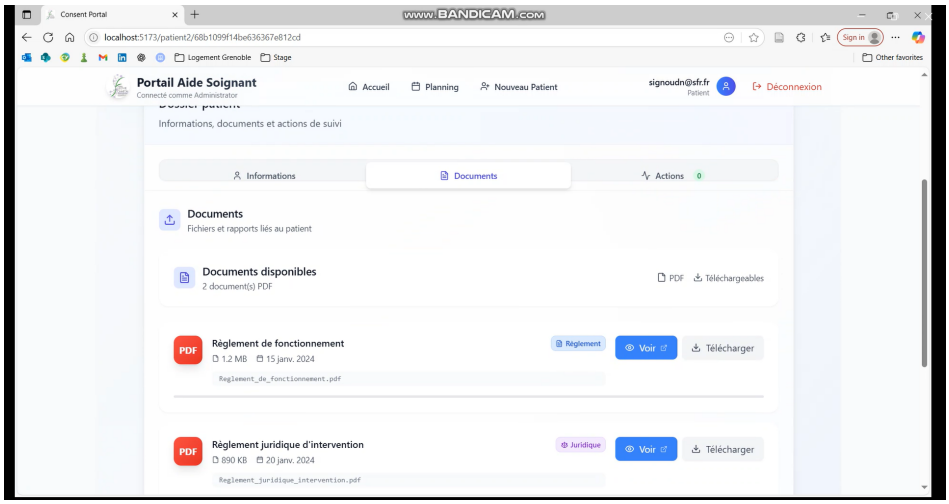


Figure A.2 – Document personalization

A.0.3 Care Planning and Calendar Module

In addition to document management, a **calendar module** was developed to support healthcare staff in organizing patient visits and interventions. The interface provides a clear view of appointments (daily, weekly, or monthly), with features such as drag-and-drop scheduling, linking visits to patient records, and logging notes after each intervention.

Future extensions of this module include integrating **route optimization** algorithms to minimize travel time for home visits. The approach under study is based on the Traveling Salesman Problem (TSP), using patient addresses geocoded into coordinates and applying heuristic graph algorithms to compute optimized visitation orders.

To ensure semantic interoperability, the system integrates the **International Classification for Nursing Practice (ICNP)** for encoding nursing interventions and outcomes. This integration aligns the project with international nursing data standards, facilitates knowledge reuse, and enhances interoperability with other e-health systems. Figure A.3 presents how the calender module implementation of the solution.

A.0.4 Security and Perspectives

The different components of the application were run locally, for both data handling and LLM generation. This respects well the RGPD guidelines dictated by the law, and the internal

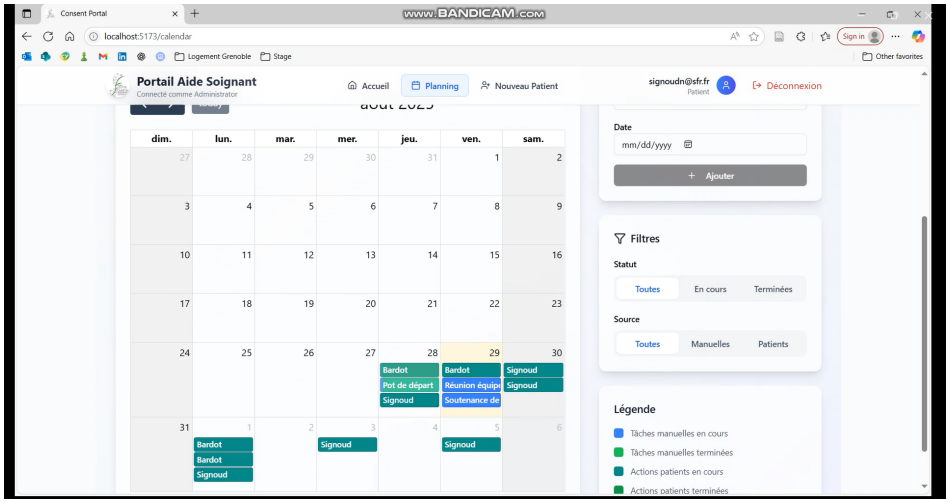


Figure A.3 – Care planning

regulations of the health institution.

The developed prototype provides an overview of our proposed approach, being displayed in real-world scenarios. The next step will consist of conducting **user evaluations with healthcare professionals** to assess usability, clinical relevance, and compliance of the generated documents. This experimental phase will provide valuable insights for refining the system and validating its potential in practical healthcare settings.

A.0.5 Conclusion

This second part of the chapter demonstrated how the proposed approach can be used within a real institutional workflow through a complete proof-of-concept system. By modeling the operational processes of a home-care institution and integrating both the personalized consent pipeline and the care recommendation engine, the prototype validates the practical feasibility of our approach.

The resulting web application shows that ontology-based reasoning, LLM-driven personalization, and multi-label care prediction can support healthcare professionals in their daily tasks. The system ensures that consent information is automatically adapted to each patient's profile, while nursing actions are recommended in accordance with standardized care terminology and institutional constraints. It also illustrates how semantic technologies can bridge heterogeneous data sources and how LLMs can be used not only for content generation but also for alignment, explanation, and interaction.

Overall, this proof of concept demonstrates the end-to-end applicability of the contributions presented in this thesis—from semantic modeling and data standardization to predictive analytics and user-facing interfaces.

Bibliography

[1] Tom L Beauchamp and James F Childress. *Principles of Biomedical Ethics*. Oxford University Press, 2019 (cit. on pp. 1, 10, 14).

[2] World Medical Association. *Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects*. 2013 (cit. on p. 1).

[3] Neil C Manson and Onora O'Neill. *Rethinking Informed Consent in Bioethics*. Cambridge University Press, 2007 (cit. on p. 1).

[4] *Nouveau guide du Conseil de l'Europe sur la littératie en santé : favoriser la confiance et l'accès équitable aux soins de santé*. Conseil de l'Europe. 2023. URL: https://www.coe.int/fr/web/human-rights-and-biomedicine/-/nouveau-guide-du-conseil-de-l-europe-sur-la-litt%C3%A9racy-en-sant%C3%A9-favoriser-la-confiance-et-l-acc%C3%A8s-%C3%A9quitable-aux-soins-de-sant%C3%A9-1 (cit. on p. 1).

[5] *Référentiel Force Probante des documents de santé — Annexe 6 : Classification des documents de santé*. Accessed: 2025-10-13. Agence du Numérique en Santé (France). 2021. URL: https://esante.gouv.fr/espace_documentation/referentiel-force-probante/annexe-6-classification-des-documents-de-sante/actual (cit. on p. 1).

[6] Lin Zhou and et al. "An ontology-based framework for informed consent in biomedical research". In: *BMC Medical Ethics* 20.1 (2019), p. 33 (cit. on p. 2).

[7] Kevin Donnelly. "SNOMED-CT: The advanced terminology and coding system for eHealth". In: *Studies in Health Technology and Informatics* 121 (2006), p. 279 (cit. on pp. 2, 52, 71, 98, 102).

[8] Tom B Brown et al. "Language Models are Few-Shot Learners". In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 1877–1901 (cit. on p. 2).

[9] Hugo Touvron et al. "LLaMA: Open and Efficient Foundation Language Models". In: *arXiv preprint arXiv:2302.13971* (2023) (cit. on p. 2).

[10] Patrick Lewis and et al. "Retrieval-augmented generation for knowledge-intensive NLP tasks". In: *Advances in Neural Information Processing Systems* 33 (2020), pp. 9459–9474 (cit. on p. 2).

[11] Jason Wei et al. "Chain-of-thought prompting elicits reasoning in large language models". In: *Proceedings of the 36th International Conference on Neural Information Processing Systems*. NIPS '22. New Orleans, LA, USA: Curran Associates Inc., 2022 (cit. on pp. 2, 26, 74).

[12] *Création d'une chaire CIREX Léonie Chaptal – CYU*. CY Cergy Paris Université. URL: https://www.cyu.fr/universite/media-et-communication/communiques-de-presse/creation-dune-chaire-cirex-leonie-chaptal-cyu (cit. on p. 2).

[13] *Fondation Léonie Chaptal*. Fondation Léonie Chaptal. URL: http://www.fondationleoniechaptal.fr/ (cit. on p. 2).

[14] Flavie SARTORI. "Le consentement libre et éclairé des personnes vulnérables". In: *LORRAINE University* (2021) (cit. on p. 10).

[15] Tomasz Pietrzykowski and Katarzyna Smilowska. "The reality of informed consent: empirical studies on patient comprehension—systematic review". In: *Trials* 22 (2021), p. 57 (cit. on p. 11).

[16] Nouri S. Fernandez A. Sudore R. L. Schillinger D. Klein-Fedyshin M. Schenker Y. Glaser J. "Interventions to Improve Patient Comprehension in Informed Consent for Medical and Surgical Procedures: An Updated Systematic Review". In: *Medical decision making : an international journal of the Society for Medical Decision Making* (2020) (cit. on p. 11).

[17] L. J. Convie, J. M. Clements, S. McCain, et al. "Development of a core outcome set for informed consent for therapy: an international key stakeholder consensus study". In: *BMC Medical Ethics* 23 (2022), p. 79 (cit. on p. 11).

[18] M. C. Plaiasu, D. O. Alexandru, and C. A. Nanu. "Physicians' legal knowledge of informed consent and confidentiality: a cross-sectional study". In: *BMC Medical Ethics* 23 (2022), p. 93 (cit. on p. 11).

[19] Díaz-Agea J. L. Suárez-Cortés M. Molina-Rodríguez A. Jiménez-Ruiz I. García-Sánchez A. García-Álvarez J. M. "Formal Quality and Compliance of Informed Consent Forms in Critical Care and Surgical Areas in Spain: An Observational Study". In: *Nursing reports* (2022) (cit. on p. 12).

[20] Claude-Audrey Meguieze and Eric Nseme Etouckey. "The informed consent in pediatric care". In: *Health Research in Africa* 3 (2025), p. 6837 (cit. on p. 12).

[21] A. G. Dunn et al. "Informed consent in clinical research in France: assessment and factors associated with therapeutic misconception". In: *Journal of Medical Ethics* (2008) (cit. on p. 12).

[22] Y. Hénin et al. "Aide à la rédaction du document écrit destiné à l'information du participant à la recherche biomédicale et à l'attestation de son consentement éclairé". In: *Thérapie* 65.2 (2010), pp. 71–94 (cit. on p. 12).

[23] Antoine Glauzy, Bernard Baertschi, and Jean-Charles Duclos-Vallée. "Le consentement médical à l'ère de la médecine de précision: enjeux et recommandations". In: *Médecine/Sciences* 39.8-9 (2023), pp. 658–663 (cit. on p. 12).

[24] Linda L. Emanuel Ezekiel J. Emanuel. "Four Models of the Physician-Patient Relationship". In: (1992) (cit. on p. 12).

[25] Bonnard Mathilde. "Le consentement du patient : du paternalisme médical à l'autonomie". Theses. Université de Lorraine, Dec. 2024 (cit. on p. 12).

[26] Shruti Jha, Chaitanya V Mahamuni, and Ishmeen Kaur Garewal. "Natural Language Processing: A Literature Survey of Approaches, Applications, Current Trends, and Future Directions". In: *2024 Asian Conference on Intelligent Technologies (ACOIT)*. 2024, pp. 1–12 (cit. on p. 15).

[27] Niki Parmar Jakob Uszkoreit Llion Jones Aidan N Gomez Łukasz Kaiser Ashish Vaswani Noam Shazeer and Illia Polosukhin. "Attention is all you need". In: *Proceedings of Advances in Neural Information Processing Systems, volume 30* (2017) (cit. on pp. 16, 22).

[28] Jürgen Schmidhuber. "Deep learning in neural networks: An overview". In: *Neural networks* (2015) (cit. on p. 16).

[29] Ruiqi Li, Xiang Zhao, and Marie-Francine Moens. "A Brief Overview of Universal Sentence Representation Methods: A Linguistic View". In: *ACM Comput. Surv.* 55.3 (Mar. 2022) (cit. on p. 16).

[30] Dan Popovici Yoshua Bengio Pascal Lamblin and Hugo Larochelle. "Greedy layerwise training of deep networks". In: *Advances in neural information processing systems* (2006) (cit. on p. 17).

[31] Yoshua Bengio Dumitru Erhan Aaron Courville and Pascal Vincent. "unsupervised pre-training help deep learning?" In: *In Proceedings of the thirteenth international conference on artificial intelligence and statistics* (2010) (cit. on p. 17).

[32] Tong Xiao and Jingbo Zhu. *Foundations of Large Language Models*. 2025. arXiv: `2501.09223 [cs.CL]` (cit. on p. 17).

[33] Improving language understanding by generative pre training. "Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever". In: *OpenAI* (2018) (cit. on p. 18).

[34] Kenton Lee Jacob Devlin Ming-Wei Chang and Kristina Toutanova. "Bert: Pretraining of deep bidirectional transformers for language understanding". In: *In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (2019) (cit. on p. 19).

[35] Adam Roberts Katherine Lee Sharan Narang Michael Matena Yanqi Zhou Wei Li Colin Raffel Noam Shazeer and Peter J. Liu. "Exploring the limits of transfer learning with a unified text-to-text transformer". In: *Journal of Machine Learning Research, 21(140)* (2020) (cit. on p. 19).

[36] Ronen Eldan Johannes Gehrke Eric Horvitz Ece Kamar Peter Lee Yin Tat Lee Yuanzhi Li-Scott M. Lundberg Harsha Nori Hamid Palangi Marco Túlio Ribeiro Sébastien Bubeck Varun Chandrasekaran and Yi Zhang. "Sparks of artificial general intelligence: Early experiments with gpt-4". In: *arXiv* (2023) (cit. on p. 19).

[37] Tom Henighan Tom B Brown Benjamin Chess Rewon Child Scott Gray Alec Radford Jeffrey Wu Jared Kaplan Sam McCandlish and Dario Amodei. "Scaling laws for neural language models". In: *arXiv* (2020) (cit. on p. 23).

[38] Victor Sanh et al. "Multitask prompted training enables zero-shot task generalization". In: *In Proceedings of International Conference on Learning Representations* (2022) (cit. on p. 23).

[39] Pranab Sahoo et al. "A Systematic Survey of Prompt Engineering in Large Language Models: Techniques and Applications". In: *ArXiv* abs/2402.07927 (2024) (cit. on p. 25).

[40] Alec Radford et al. "Language Models are Unsupervised Multitask Learners". In: 2019 (cit. on p. 25).

[41] Tom Brown et al. “Language Models are Few-Shot Learners”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle et al. Vol. 33. Curran Associates, Inc., 2020, pp. 1877–1901 (cit. on p. 25).

[42] Zhuosheng Zhang et al. “Automatic Chain of Thought Prompting in Large Language Models”. In: *ArXiv* abs/2210.03493 (2022) (cit. on p. 26).

[43] Xuezhi Wang et al. “Self-Consistency Improves Chain of Thought Reasoning in Language Models”. In: *ArXiv* abs/2203.11171 (2022) (cit. on p. 26).

[44] Hanmeng Liu et al. “LogiCoT: Logical Chain-of-Thought Instruction Tuning”. In: *The 2023 Conference on Empirical Methods in Natural Language Processing*. 2023 (cit. on p. 27).

[45] Hanxu Hu et al. *Chain-of-Symbol Prompting Elicits Planning in Large Langauge Models*. 2024. arXiv: 2305.10276 [cs.CL] (cit. on p. 27).

[46] Shunyu Yao et al. “Tree of Thoughts: Deliberate Problem Solving with Large Language Models”. In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023 (cit. on p. 28).

[47] Maciej Besta et al. “Graph of thoughts: solving elaborate problems with large language models”. In: *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence*. AAAI Press, 2024 (cit. on p. 28).

[48] Jason E. Weston and Sainbayar Sukhbaatar. “System 2 Attention (is something you might need too)”. In: *ArXiv* abs/2311.11829 (2023) (cit. on p. 28).

[49] Yucheng Zhou et al. “Thread of Thought Unraveling Chaotic Contexts”. In: *ArXiv* abs/2311.08734 (2023) (cit. on p. 29).

[50] Zilong Wang et al. “Chain-of-Table: Evolving Tables in the Reasoning Chain for Table Understanding”. In: *The Twelfth International Conference on Learning Representations*. 2024 (cit. on p. 29).

[51] Hassan S. Al Khatib et al. “Patient-centric knowledge graphs: a survey of current methods, challenges, and applications”. In: *Frontiers in Artificial Intelligence* Volume 7 - 2024 (2024) (cit. on p. 30).

[52] Hejie Cui et al. “A Survey on Knowledge Graphs for Healthcare: Resources, Application Progress, and Promise”. In: *ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH)*. 2023 (cit. on p. 30).

[53] Holger Wache et al. “Ontology-Based Integration of Information-A Survey of Existing Approaches.” In: *Ois@ ijcai*. 2001 (cit. on pp. 30, 31).

[54] Mohamed Hedi Karray, Brigitte Chebel-Morello, and Noureddine Zerhouni. “A formal ontology for industrial maintenance”. In: *Applied ontology* 7.3 (2012), pp. 269–310 (cit. on p. 31).

[55] Franz Baader et al., eds. *The Description Logic Handbook: Theory, Implementation and Applications*. 2nd ed. Cambridge: Cambridge University Press, 2007 (cit. on pp. 31, 98).

[56] Richong Zhang, Yongyi Mao, and Weihua Zhao. “Knowledge Graphs Completion via Probabilistic Reasoning”. In: *Information Sciences* 521 (2020), pp. 144–159 (cit. on p. 31).

[57] Dagobert Soergel. “Digital Libraries and Knowledge Organization”. In: *Semantic Digital Libraries*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 9–39 (cit. on pp. 31, 32).

[58] Marlies Olensky. “Semantic interoperability in Europeana: An Examination of CIDOC CRM in Digital Cultural Heritage Documentation”. In: *Bulletin of IEEE Technical Committee on Digital Libraries 6 (2)* (2010) (cit. on p. 31).

[59] Dragan Gasevic, Dragan Djuric, and Vladan Devedzic. *Model Driven Engineering and Ontology Development*. 2nd. Springer Publishing Company, Incorporated, 2009 (cit. on p. 32).

[60] Matteo Cristani and Roberta Cuel. “A Survey on Ontology Creation Methodologies”. In: *Int. J. Semantic Web Inf. Syst.* 1.2 (2005), pp. 49–69 (cit. on p. 32).

[61] Jérôme Euzenat and Pavel Shvaiko. *Ontology Matching*. Springer, 2013 (cit. on p. 33).

[62] Enrico G. Caldarola, Antonio Picariello, and Antonio M. Rinaldi. “An approach to ontology integration for ontology reuse in knowledge based digital ecosystems”. In: New York, NY, USA: Association for Computing Machinery, 2015 (cit. on p. 33).

[63] Sirko Schindler and Jan Martin Keil. “Building Ontologies for Reuse: Lessons Learned from Unit Ontologies”. In: *Joint Ontology Workshops*. 2019 (cit. on p. 33).

[64] Nikolaos Trokanas. “Practical issues of ontology design, evaluation and reuse: a case study in chemical and process engineering”. In: 2016 (cit. on p. 33).

[65] Alessia Pisu et al. “Leveraging Language Models for Generating Ontologies of Research Topics”. In: *TEXT2KG/DQMLKG@ESWC*. 2024 (cit. on p. 33).

[66] Maurice Funk et al. “Towards Ontology Construction with Language Models”. In: *ArXiv* abs/2309.09898 (2023) (cit. on p. 33).

[67] Tim Berners-Lee, JAMES Hendler, and Ora Lassila. “A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities”. In: *Scientific American* (2002) (cit. on p. 33).

[68] Tim Berners-Lee, James Hendler, and Ora Lassila. “The semantic web”. In: *Scientific american* 284.5 (2001), pp. 34–43 (cit. on pp. 33, 34).

[69] RDF. *Resource Description Framework*. URL: `https://www.w3.org/RDF/`. 2014 (cit. on p. 34).

[70] Grigoris Antoniou and Frank Van Harmelen. “Web ontology language: Owl”. In: *Handbook on ontologies*. Springer, 2004, pp. 67–92 (cit. on p. 35).

[71] David Allemang and James Hendler. *Semantic web for the working ontologist: effective modeling in RDFS and OWL*. Morgan Kaufmann, 2011 (cit. on p. 35).

[72] Y. Guo, R. Fikes, and F. van Harmelen. “OWL Web Ontology Language Overview”. In: *W3C Recommendation*. 2004 (cit. on pp. 35, 36).

[73] Pascal Hitzler, Markus Krötzsch, and Sebastian Rudolph. *Foundations of Semantic Web Technologies*. CRC Press, 2009 (cit. on pp. 35, 36).

[74] W3C RDF Working Group. *RDF 1.1 Concepts and Abstract Syntax*. 2014 (cit. on p. 35).

[75] W3C RDF Working Group. *RDF Schema 1.1*. 2014 (cit. on p. 35).

[76] W3C OWL Working Group. *OWL 2 Web Ontology Language Document Overview (Second Edition)*. 2012 (cit. on p. 36).

[77] Eric Prud'Hommeaux, Andy Seaborne, et al. "Sparql query language for rdf (working draft)". In: *Bericht, W3C, March* (2007) (cit. on p. 36).

[78] Ian Horrocks et al. "SWRL: A semantic web rule language combining OWL and RuleML". In: *W3C Member submission* 21.79 (2004), pp. 1–31 (cit. on p. 37).

[79] Bhaskar Kapoor and Savita Sharma. "A comparative study ontology building tools for semantic web applications". In: *International journal of Web & Semantic Technology (IJWesT)* 1.3 (2010), pp. 1–13 (cit. on p. 37).

[80] Protégé. *"Stanford University"*. URL: `http://protege.stanford.edu/`. 2014 (cit. on p. 37).

[81] Birte Glimm et al. "HermiT: an OWL 2 reasoner". In: *Journal of Automated Reasoning* 53.3 (2014), pp. 245–269 (cit. on p. 37).

[82] Bijan Parsia and Evren Sirin. "Pellet: An owl dl reasoner". In: *Third international semantic web conference-poster*. Vol. 18. Publishing. 2004, p. 13 (cit. on p. 37).

[83] Matthew Horridge and Sean Bechhofer. "The owl api: A java api for owl ontologies". In: *Semantic web* 2.1 (2011), pp. 11–21 (cit. on p. 37).

[84] Jean-Baptiste Lamy. "Owlready: Ontology-oriented programming in Python with automatic classification and high level constructs for biomedical ontologies". In: *Artificial intelligence in medicine* 80 (2017), pp. 11–28 (cit. on p. 37).

[85] *QUÉBEC, Loi sur les services de santé et les services sociaux: L.R.Q., chapitre S-4.2*. 2011 (cit. on p. 45).

[86] X. Wu and et al. "Medical Knowledge Graph: Data Sources, Construction, Reasoning, and Applications". In: *Big Data Mining and Analytics* 6.2 (2023), pp. 201–217 (cit. on pp. 45, 70).

[87] Wisam A. Qader, Musa M. Ameen, and Bilal I. Ahmed. "An Overview of Bag of Words;Importance, Implementation, Applications, and Challenges". In: *2019 International Engineering Conference (IEC)*. 2019, pp. 200–204 (cit. on pp. 50, 53).

[88] Sarika Zaware et al. "Text Summarization using TF-IDF and Textrank algorithm". In: *2021 5th International Conference on Trends in Electronics and Informatics (ICOEI)*. 2021, pp. 1399–1407 (cit. on pp. 50, 53).

[89] Wang and et al. "Measurement of Text Similarity: A Survey". In: *Information* (Sept. 2020) (cit. on p. 50).

[90] Hai Zhou. "Research of Text Classification Based on TF-IDF and CNN-LSTM". In: *Journal of Physics: Conference Series* (Jan. 1, 2022) (cit. on p. 50).

[91] Joyinee Dasgupta et al. “A Survey of Numerous Text Similarity Approach”. In: *International Journal of Scientific Research in Computer Science, Engineering and Information Technology* (Feb. 22, 2023) (cit. on p. 50).

[92] Wael H.Gomaa and Aly A. Fahmy. “A Survey of Text Similarity Approaches”. In: *International Journal of Computer Applications* (Apr. 18, 2013) (cit. on pp. 51, 53).

[93] Elvys Linhares Pontes et al. *Predicting the Semantic Textual Similarity with Siamese CNN and LSTM*. Oct. 24, 2018 (cit. on pp. 51, 53).

[94] Md T. R. Laskar and et al. “Contextualized Embeddings based Transformer Encoder for Sentence Similarity Modeling in Answer Selection Task”. In: *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, May 2020 (cit. on pp. 51, 53).

[95] Sunwoo Ahn et al. “Practical Binary Code Similarity Detection with BERT-based Transferable Similarity Learning”. In: *Proceedings of the 38th Annual Computer Security Applications Conference*. New York, NY, USA: Association for Computing Machinery, Dec. 5, 2022 (cit. on pp. 51, 53).

[96] Chandra Yadav and Aditi Sharan. “A New LSA and Entropy-Based Approach for Automatic Text Document Summarization”. In: *International Journal on Semantic Web and Information Systems (IJSWIS)* (Oct. 1, 2018) (cit. on pp. 51, 53).

[97] Mingyu Ji and Xinhai Zhang. “A Short Text Similarity Calculation Method Combining Semantic and Headword Attention Mechanism”. In: *Scientific Programming* (May 21, 2022) (cit. on pp. 51, 53).

[98] Somaiyeh Dehghan and Mehmet Fatih Amasyali. “SupMPN: Supervised Multiple Positives and Negatives Contrastive Learning Model for Semantic Textual Similarity”. In: *Applied Sciences* (Jan. 2022) (cit. on pp. 51, 53).

[99] F. Ahmad and M. Faisal. “A novel hybrid methodology for computing semantic similarity between sentences through various word senses”. In: *International Journal of Cognitive Computing in Engineering* (June 1, 2022) (cit. on pp. 51, 53).

[100] Z. Amur et al. “Short-Text Semantic Similarity (STSS): Techniques, Challenges and Future Perspectives”. In: *Applied Sciences* (Jan. 2023) (cit. on pp. 51, 53).

[101] World Health Organization. *International Classification of Diseases 11th Revision (ICD-11)*. World Health Organization, 2019 (cit. on p. 52).

[102] World Health Organization. *International Classification of Functioning, Disability and Health (ICF)*. World Health Organization, 2001 (cit. on p. 52).

[103] Olivier Bodenreider. “The Unified Medical Language System (UMLS): integrating biomedical terminology”. In: *Nucleic Acids Research* 32.suppl_1 (2004), pp. D267–D270 (cit. on p. 52).

[104] Alexandru Constantin, Silvio Peroni, and David Shotton. “The Document Components Ontology (DoCO)”. In: *Semantic Web Journal* 5.2 (2014), pp. 101–107 (cit. on p. 52).

[105] Chryssoula Bekiari et al. “FRBRoo: Enabling a Common View of Information from Memory Institutions”. In: *Digital Libraries: International Collaboration and Cross-Fertilization*. Vol. 10020. Lecture Notes in Computer Science. Springer, 2016, pp. 30–40 (cit. on p. 54).

[106] Martin Doerr. “The CIDOC Conceptual Reference Model: An Ontological Approach to Semantic Interoperability of Metadata”. In: *AI Magazine*. Vol. 24. 3. 2003, pp. 75–92 (cit. on p. 54).

[107] Clement J. McDonald et al. “LOINC, a universal standard for identifying laboratory observations: a 5-year update”. In: *Clinical Chemistry* 49.4 (2003), pp. 624–633 (cit. on p. 54).

[108] S. C. Grant. “Informed Consent—We Can and Should Do Better”. In: *JAMA Netw Open* (2021) (cit. on p. 69).

[109] S. L. Fleming and A. et al. Lozano. “MedAlign: A Clinician-Generated Dataset for Instruction Following with Electronic Medical”. In: *arXiv preprint arXiv:2308.14089* (2023) (cit. on p. 69).

[110] C. Sinsky and et al. “Allocation of physician time in ambulatory practice: a time and motion study in 4 specialties”. In: *Annals of Internal Medicine* 165 (2016), pp. 753–760 (cit. on p. 69).

[111] K. E. Robinson and J. A. Kersey. “Novel electronic health record (EHR) education intervention in large healthcare organization improves quality, efficiency, time, and impact on burnout”. In: *Medicine* (2018) (cit. on p. 69).

[112] A. J. Thirunavukarasu et al. “Large language models in medicine”. In: *Nature Medicine* (2023) (cit. on p. 69).

[113] *Le référentiel de bonnes pratiques professionnelles pour le service de soins infirmiers à domicile, 2018-2023*. Accessed via: `https://www.normandie.ars.sante.fr/media/35329/download?inline`. 2018 (cit. on p. 69).

[114] Tao Tu et al. *Towards Generalist Biomedical AI*. 2023. arXiv: `2307.14334 [cs.CL]` (cit. on p. 69).

[115] Jinhyuk Lee and Wonjin et al. Yoon. “BioBERT: a pre-trained biomedical language representation model for biomedical text mining”. In: *Bioinformatics* (2019) (cit. on p. 69).

[116] Kexin Huang, Jaan Altosaar, and Rajesh Ranganath. “ClinicalBERT: Modeling Clinical Notes and Predicting Hospital Readmission”. In: *ArXiv* abs/1904.05342 (2019) (cit. on p. 69).

[117] Chaitanya Shivade. “MedNLI - A Natural Language Inference Dataset For The Clinical Domain”. In: *PhysioNet* (2019) (cit. on p. 69).

[118] Z. Hu et al. “Empowering language models with knowledge graph reasoning for open-domain question answering”. In: *In EMNLP* (2022) (cit. on p. 69).

[119] X. Zhang et al. “Greaselm: Graph reasoning enhanced language models”. In: *In ICLR* (2022) (cit. on p. 69).

[120] Urvashi Khandelwal et al. “Generalization through memorization: Nearest neighbor language models”. In: *arXiv preprint arXiv:1911.00172* (2019) (cit. on p. 70).

[121] Sosuke Nishikawa et al. “EASE: Entity-Aware Contrastive Learning of Sentence Embedding”. In: *North American Chapter of the Association for Computational Linguistics*. 2022 (cit. on p. 70).

[122] D.M. Anisuzzaman et al. “Fine-Tuning Large Language Models for Specialized Use Cases”. In: *Mayo Clinic Proceedings: Digital Health* 3.1 (2025), p. 100184 (cit. on p. 70).

[123] Yunfan Gao et al. *Retrieval-Augmented Generation for Large Language Models: A Survey*. 2024. arXiv: `2312.10997 [cs.CL]` (cit. on p. 70).

[124] Diego Sanmartin. *KG-RAG: Bridging the Gap Between Knowledge and Creativity*. 2024. arXiv: `2405.12035 [cs.AI]` (cit. on pp. 70, 74).

[125] Chengrui Wang et al. *BioRAG: A RAG-LLM Framework for Biological Question Reasoning*. 2024. arXiv: `2408.01107 [cs.CL]` (cit. on pp. 70, 74).

[126] Haochen Liu et al. “Knowledge Graph-Enhanced Large Language Models via Path Selection”. In: *ArXiv* abs/2406.13862 (2024) (cit. on pp. 70, 74).

[127] Shilong Li et al. *GraphReader: Building Graph-based Agent to Enhance Long-Context Abilities of Large Language Models*. 2024. arXiv: `2406.14550 [cs.CL]` (cit. on pp. 70, 74).

[128] Junjie Wang et al. “Learning to Plan for Retrieval-Augmented Large Language Models from Knowledge Graphs”. In: *Findings of the Association for Computational Linguistics: EMNLP 2024*. Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 7813–7835 (cit. on pp. 70, 74).

[129] Ankan Mullick et al. “On The Persona-based Summarization of Domain-Specific Documents”. In: *Findings of the Association for Computational Linguistics: ACL 2024*. Ed. by Lun-Wei Ku, Andre Martins, and Vivek Srikumar. Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 14291–14307 (cit. on pp. 70, 86, 90).

[130] Tiannan Wang et al. *AI PERSONA: Towards Life-long Personalization of LLMs*. 2024. arXiv: `2412.13103 [cs.CL]` (cit. on p. 70).

[131] Liana Razmerita and Miltiadis D. Lytras. “Ontology-Based User Modelling Personalization: Analyzing the Requirements of a Semantic Learning Portal”. In: *Emerging Technologies and Information Systems for the Knowledge Society*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 354–363 (cit. on p. 71).

[132] Alberti V et al. “Ontology engineering to model the European cultural heritage: The case of Cultural gems”. In: *Lecture Notes in Networks and Systems* 696 (2024), pp. 1–10 (cit. on p. 71).

[133] Esther Román-Villarán et al. “A Personalized Ontology-Based Decision Support System for Complex Chronic Patients: Retrospective Observational Study”. In: *JMIR Formative Research* 6.8 (2022) (cit. on p. 71).

[134] Matilde Pato, Márcia Barros, and Francisco M. Couto. “Survey on Recommender Systems for Biomedical Items in Life and Health Sciences”. In: *ACM Comput. Surv.* 56.6 (Feb. 2024) (cit. on p. 71).

[135] Dan Brickley. “FOAF Vocabulary Specification”. In: 2004 (cit. on pp. 71, 98).

[136] Duane Bender and Kamran Sartipi. “HL7 FHIR: An Agile and RESTful approach to healthcare information exchange”. In: *Proceedings of the 26th IEEE International Symposium on Computer-Based Medical Systems*. 2013, pp. 326–331 (cit. on pp. 71, 98).

[137] *Enhanced Personal Profile Ontology | NCBO BioPortal : E-PPO* (cit. on pp. 71, 98).

[138] *Clinical Smell and Taste Disorders | NCBO BioPortal : CSTD* (cit. on pp. 71, 72, 98).

[139] *Health and Social Person-centric Ontology | NCBO BioPortal : HSPO* (cit. on pp. 71, 98).

[140] *Culture, Race, Ethnicity and Nationality Ontology | NCBO BioPortal : CRENO* (cit. on pp. 71, 98).

[141] Nikit Srivastava et al. *MST5 – Multilingual Question Answering over Knowledge Graphs*. 2024. arXiv: 2407.06041 [cs.CL] (cit. on p. 74).

[142] Xingxuan Li et al. “Chain-of-Knowledge: Grounding Large Language Models via Dynamic Knowledge Adapting over Heterogeneous Sources”. In: *International Conference on Learning Representations*. 2023 (cit. on p. 74).

[143] Yifu Gao et al. “Two-stage Generative Question Answering on Temporal Knowledge Graph Using Large Language Models”. In: *Annual Meeting of the Association for Computational Linguistics*. 2024 (cit. on p. 74).

[144] Dawei Li et al. “DALK: Dynamic Co-Augmentation of LLMs and KG to answer Alzheimer’s Disease Questions with Scientific Literature”. In: *Findings of the Association for Computational Linguistics: EMNLP 2024*. Ed. by Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen. Miami, Florida, USA: Association for Computational Linguistics, Nov. 2024, pp. 2187–2205 (cit. on p. 74).

[145] Hongliu Cao et al. *Multi-Agent LLM Judge: automatic personalized LLM judge design for evaluating natural language generation applications*. 2025. arXiv: 2504.02867 [cs.CL] (cit. on p. 90).

[146] Jiongnan Liu et al. “LLMs + Persona-Plug = Personalized LLMs”. In: *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Wanxiang Che et al. Vienna, Austria: Association for Computational Linguistics, July 2025, pp. 9373–9385 (cit. on p. 90).

[147] Amira Henaien, Hadda Ben Elhadj, and Lamia Chaari Fourati. “Combined Machine Learning and Semantic Modelling for Situation Awareness and Healthcare Decision Support”. In: *The Impact of Digital Technologies on Public Health in Developed and Developing Countries*. Cham: Springer International Publishing, 2020, pp. 197–209 (cit. on p. 97).

[148] Xuehong Wu et al. “Medical Knowledge Graph: Data Sources, Construction, Reasoning, and Applications”. In: *Big Data Mining and Analytics* 6.2 (June 2023), pp. 201–217 (cit. on p. 97).

[149] V. K. C. Yan, X. Li, and X. Ye. *Drug Repurposing for the Treatment of COVID-19: A Knowledge Graph Approach*. 2021 (cit. on p. 97).

[150] Xiang Wang et al. “KGAT: Knowledge Graph Attention Network for Recommendation”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. KDD ’19. New York, NY, USA: Association for Computing Machinery, July 2019, pp. 950–958 (cit. on pp. 97, 98).

[151] Silvia Schiaffino and Analía Amandi. “Building an Expert Travel Agent as a Software Agent”. In: *Expert Systems with Applications* 36.2, Part 1 (Mar. 2009), pp. 1291–1299 (cit. on p. 98).

[152] Fuming Sun et al. “Multi-Label Image Categorization with Sparse Factor Representation”. In: *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society* 23.3 (Mar. 2014), pp. 1028–1037 (cit. on p. 98).

[153] Anastasios Dimou et al. “An Empirical Study of Multi-label Learning Methods for Video Annotation”. In: *2009 Seventh International Workshop on Content-Based Multimedia Indexing*. 2009, pp. 19–24 (cit. on p. 98).

[154] Francisco Herrera et al. *Multilabel Classification*. Cham: Springer International Publishing, 2016 (cit. on p. 98).

[155] Thuan Pham et al. “Graph-Based Multi-Label Disease Prediction Model Learning from Medical Data and Domain Knowledge”. In: *Knowledge-Based Systems* 235 (Jan. 2022), p. 107662 (cit. on p. 98).

[156] Joyce Zhou, Yijia Dai, and Thorsten Joachims. *Language-Based User Profiles for Recommendation*. 2024. arXiv: `2402.15623 [cs.CL]` (cit. on p. 98).

[157] Fan Yang et al. *PALR: Personalization Aware LLMs for Recommendation*. 2023. eprint: `2305.07622` (cs) (cit. on p. 98).

[158] “A survey on augmenting knowledge graphs (KGs) with large language models (LLMs): models, evaluation metrics, benchmarks, and challenges”. In: *Discov Artif Intell* (2024) (cit. on p. 98).

[159] *ICNP Browser*. https://www.icn.ch/icnp-browser (cit. on p. 101).

Résumé — Cette thèse porte sur la personnalisation du consentement éclairé dans le domaine de la santé. Elle cherche à répondre à la question centrale : Comment concevoir un système capable de produire un consentement médical qui soit à la fois conforme à la loi, compréhensible et adapté au profil du patient, et assistant le personnel soignant à la proposition des soins ? Pour répondre à cette problématique, la thèse s'articule autour de trois piliers principaux : la conformité légale, la personnalisation du consentement, et la recommandation de soins. La solution proposée explore l'intersection entre les technologies sémantiques et l'intelligence artificielle générative, à travers le système SPIRIT (Semantic Personalized and Compliant System for Informed consent and care Recommendation). Les principales contributions de ce travail comprennent : (i) une approche basée sur l'ontologie pour garantir la conformité légale, (ii) un cadre de raisonnement sémantique pour la génération de consentements personnalisés, et (iii) un système de recommandation infirmière fondé sur des graphes de connaissances. Enfin, une preuve de concept fonctionnelle valide la faisabilité de l'approche et ouvre des perspectives d'expérimentation en conditions réelles, notamment vers des extensions multilingues et multimodales.

Mots clés : Consentement éclairé, Conformité légale, Personnalisation, Graphe de connaissances, Ontologie, LLM, NLP, Recommandation.

Abstract — This thesis focuses on the personalization of informed consent in healthcare. It seeks to address the central question: How can we design a system capable of producing a medical consent that is simultaneously legally compliant, comprehensible, and adapted to the patient's profile, while supporting healthcare professionals in proposing appropriate care plans? To answer this question, the research is structured around three key pillars: legal compliance, consent personalization, and care recommendation. The proposed solution explores the intersection of semantic technologies and generative AI, by proposing SPIRIT (Semantic Personalized and Compliant System for Informed consent and care Recommendation). The main contributions of this work include: (i) an ontology-driven approach for ensuring legal compliance, (ii) a semantic reasoning framework for personalized consent generation, and (iii) a knowledge-graph-based system for nursing recommendations. Finally, a functional proof of concept validates the feasibility of the approach and opens perspectives for real-world deployment, including multilingual and multimodal extensions.

Keywords: Informed Consent, Legal Compliance, Personalization, Knowledge Graph, Ontology, LLMs, NLP, Recommendation.

ETIS Laboratory UMR 8051 CY Cergy Paris University, ENSEA, CNRS, F-95000 Cergy, France