

AIX-MARSEILLE UNIVERSITÉ

ED 184 – MATHÉMATIQUES ET INFORMATIQUE
LABORATOIRE D'INFORMATIQUE FONDAMENTALE

UNIV. FEDERAL DO RIO GRANDE DO SUL
INSTITUTO DE INFORMÁTICA

Thèse présentée pour obtenir le grade universitaire de docteur

Discipline : Informatique

Silvio Ricardo CORDEIRO

Distributional models of
multiword expression compositionality prediction

Modèles distributionnels pour la
prédiction de compositionnalité d'expressions polylexicales

Soutenue le 18/12/2017 devant le jury composé de :

Timothy BALDWIN	University of Melbourne	Rapporteur
Stefan EVERT	Friedrich-Alexander-Universität	Rapporteur
Antoine DOUCET	Université de La Rochelle	Examinateur
Helena CASELI	Universidade Federal de São Carlos	Examinatrice
Aline VILLAVICENCIO	Univ. Federal do Rio Grande do Sul	Directrice de thèse
Alexis NASR	Aix-Marseille Université	Directeur de thèse
Carlos RAMISCH	Aix-Marseille Université	Directeur de thèse

Numéro national de thèse/suffixe local : 2017AIXM0501/031ED184



Cette oeuvre est mise à disposition selon les termes de la [Licence Creative Commons Attribution – Pas d’Utilisation Commerciale – Pas de Modification 4.0 International](#).

Abstract

Natural language processing systems often rely on the idea that language is compositional, that is, the meaning of a linguistic entity can be inferred from the meaning of its parts. This expectation fails in the case of multiword expressions (MWEs). For example, a person who is a *sitting duck* is neither a duck nor necessarily sitting. Modern computational techniques for inferring word meaning based on the distribution of words in the text have been quite successful at multiple tasks, especially since the rise of word embedding approaches. However, the representation of MWEs still remains an open problem in the field. In particular, it is unclear how one could predict from corpora whether a given MWE should be treated as an indivisible unit (e.g. *nut case*) or as some combination of the meaning of its parts (e.g. *engine room*). This thesis proposes a framework of MWE compositionality prediction based on representations of distributional semantics, which we instantiate under a variety of parameters. We present a thorough evaluation of the impact of these parameters on three new datasets of MWE compositionality, encompassing English, French and Portuguese MWEs. Finally, we present an extrinsic evaluation of the predicted levels of MWE compositionality on the task of MWE identification. Our results suggest that the proper choice of distributional model and corpus parameters can produce compositionality predictions that are comparable to the state of the art.

Keywords: distributional semantics, multiword expressions, compositionality, idiomaticity.

Résumé

Les systèmes de traitement automatique des langues reposent souvent sur l'idée que le langage est compositionnel, c'est-à-dire que le sens d'une entité linguistique peut être déduite à partir du sens de ses parties. Cette supposition ne s'avère pas vraie dans le cas des expressions polylexicales (EPLs). Par exemple, une personne qui est une *poule mouillée* n'est ni une poule, ni nécessairement mouillée. Les techniques modernes de calcul pour déduire le sens des mots en fonction de leur distribution dans le texte ont obtenu de bons résultats sur plusieurs tâches, en particulier depuis l'apparition des modèles communément appelés *word embeddings*. Cependant, la représentation des EPLs reste toujours un problème non résolu. En particulier, on ne sait pas comment prédire avec précision, à partir des corpus, si une EPL donnée doit être traitée comme une unité indivisible (par exemple *carton plein*) ou comme une combinaison du sens de ses parties (par exemple *eau potable*). Cette thèse propose un cadre méthodologique pour la prédiction de compositionnalité d'EPLs fondé sur des représentations de la sémantique distributionnelle, que nous avons instancié à partir d'une variété de paramètres. Nous présenterons une évaluation complète de l'impact de ces paramètres sur trois nouveaux ensembles de données modélisant la compositionnalité d'EPLs, en anglais, français et portugais. Finalement, nous présenterons une évaluation extrinsèque des niveaux de compositionnalité prédits par le modèle dans le contexte d'un système d'identification d'EPLs. Les résultats suggèrent que le choix spécifique de modèle distributionnel et de paramètres de corpus peut produire des prédictions de compositionnalité qui sont comparables à celles présentées dans l'état de l'art.

Mots clés : expressions polylexicales, sémantique distributionnelle, compositionnalité, idiomaticité.

Resumo

Sistemas de processamento de linguagem natural baseiam-se com frequência na hipótese de que a linguagem humana é composicional, ou seja, que o significado de uma entidade linguística pode ser inferido a partir do significado de suas partes. Essa expectativa falha no caso de expressões multipalavras (EMPs). Por exemplo, uma pessoa caracterizada como *pão-duro* não é literalmente um pão, e também não tem uma consistência molecular mais dura que a de outras pessoas. Técnicas computacionais modernas para inferir o significado das palavras com base na sua distribuição no texto vêm obtendo um considerável sucesso em múltiplas tarefas, especialmente após o surgimento de abordagens de *word embeddings*. No entanto, a representação de EMPs continua a ser um problema em aberto na área. Em particular, não existe um método consolidado que prediga, com base em *corpora*, se uma determinada EMP deveria ser tratada como unidade indivisível (por exemplo *olho gordo*) ou como alguma combinação do significado de suas partes (por exemplo *tartaruga marinha*). Esta tese propõe um modelo de predição de composicionalidade de EMPs com base em representações de semântica distribucional, que são instanciadas no contexto de uma variedade de parâmetros. Também é apresentada uma avaliação minuciosa do impacto desses parâmetros em três novos conjuntos de dados que modelam a composicionalidade de EMP, abrangendo EMPs em inglês, francês e português. Por fim, é apresentada uma avaliação extrínseca dos níveis previstos de composicionalidade de EMPs, através da tarefa de identificação de EMPs. Os resultados obtidos sugerem que a escolha adequada do modelo distribucional e de parâmetros de *corpus* pode produzir predições de composicionalidade que são comparáveis às observadas no estado da arte.

Palavras-chave: expressões multipalavras, semântica distribucional, composicionalidade, idiomaticidade.

Acknowledgments

My deep gratitude goes first to my advisers, Aline Villavicencio, Alexis Nasr and Carlos Ramisch, who expertly guided me through my graduate education and who shared the excitement of my three years of discovery and in-depth research in the field of NLP. Their unwavering enthusiasm for the field has kept me constantly engaged with my research, and I cannot thank them enough for these years of scientific mentorship.

I would also like to thank the reviewers of my papers, who have immensely contributed to this thesis through their feedback and advice from an external point of view. In particular, I thank the reviewers of this thesis, Timothy Baldwin and Stefan Evert, as well as the two other members of the thesis committee, Antoine Doucet and Helena Caseli.

I would especially like to thank all the coauthors of my papers. In particular, I thank my advisers, Leonardo Zilio, Rodrigo Wilkens, Marco Idiart, Manon Scholivet, Marie Candito, and Agata Savary. I have immensely enjoyed the multiple discussions we have had on how to best pursue the research at hand. I also feel the need to thank Agata Savary one more time. After all, the PARSEME network of researchers might not have existed without her. I am also indebted to the members of PARSEME, with whom I have shared research ideas in multiple occasions.

This thesis would not have been possible without the financial support from the Brazilian PhD scholarship from CAPES, along with the Brazilian project “Simplificação Textual de Expressões Complexas” sponsored by Samsung Eletrônica da Amazônia Ltda, the European project PARSEME (COST IC1207), the French project PARSEME-FR (ANR-14-CERA-0001), as well as the French-Brazilian projects CAMELEON (CAPES-COFECUB #707-11) and AIM-WEST (FAPERGS-INRIA 1706-2551/13-7).

I would also like to thank my family and friends, who have supported me through these years of PhD studies. A special thanks goes to Robert and Tanya, for their companionship in the hours when the lab lights were off.

Finally, I would like to express my gratitude to everyone who, directly or indirectly, contributed to the completion of this thesis.

Contents

Abstract	3
Résumé	4
Resumo	5
Acknowledgments	6
List of Figures	10
List of Tables	11
Abbreviations	13
1 Introduction	14
1.1 Contributions	16
1.2 Investigated hypotheses	18
1.3 Context of the thesis	21
1.4 Publications	21
1.5 Thesis structure	23
2 Background	24
2.1 Basic terminology	24
2.2 Language and statistics	27
2.2.1 Occurrence counts	28
2.2.2 Co-occurrence counts	29
2.2.3 Association measures	31
2.2.4 Inter-rater agreement measures	33
2.2.5 Evaluation measures for continuous data	34
2.2.6 Evaluation measures for categorical data	36
2.3 Multiword expressions	38
2.3.1 Nominal compounds	40
2.3.2 Type discovery	40
2.3.3 Token identification	41
2.4 Word semantics	43
2.4.1 Symbolic representations	43
2.4.2 Numerical representations	44
2.5 MWE semantics	48
2.5.1 Symbolic representation	48
2.5.2 Numerical representation	49

2.5.3	Compositionality datasets in the literature	50
3	Compositionality datasets	53
3.1	Data collection	53
3.1.1	Compound selection	54
3.1.2	Sentence selection	54
3.1.3	Questionnaire design	55
3.1.4	Judgment collection	56
3.2	Dataset analysis	57
3.2.1	Score distribution	57
3.2.2	Difficulty of annotation	57
3.2.3	Estimating whole-compound from head/modifier	59
3.2.4	Correlation with distributional variables	60
3.3	Summary	62
4	Compositionality prediction	64
4.1	Related work	65
4.2	Proposed model	67
4.3	Corpus preprocessing	71
4.4	DSMs	72
4.5	Parameters	74
4.6	Evaluation setup	75
5	Intrinsic evaluation of compositionality prediction	77
5.1	Overall highest results per DSM	77
5.1.1	English	77
5.1.2	French	79
5.1.3	Portuguese	80
5.1.4	Cross-language analysis	80
5.2	DSM parameters	81
5.2.1	Context-window size	81
5.2.2	Number of dimensions	82
5.3	Corpus parameters	83
5.3.1	Type of preprocessing	84
5.3.2	Corpus size	85
5.3.3	Parallel predictions	86
5.4	Prediction strategy	89
5.5	Sanity checks	96
5.5.1	Number of iterations	97
5.5.2	Minimum count threshold	97
5.5.3	Windows of size 2+2	98
5.5.4	Higher number of dimensions	99
5.5.5	Random initialization	99

5.5.6	Data filtering	100
5.6	Error analysis	103
5.6.1	Frequency and compositionality prediction	103
5.6.2	Conventionalization and compositionality prediction	105
5.6.3	Human–system comparison	106
5.6.4	Range-based analyses	107
5.7	Summary and discussion	110
6	Extrinsic evaluation of compositionality prediction	115
6.1	Proposed models of MWE identification	115
6.1.1	Rule-based identification model	116
6.1.2	Probabilistic identification model	117
6.2	Experimental setup	118
6.2.1	Reference corpora	118
6.2.2	Compositionality lexicons	119
6.2.3	Rule-based identification	119
6.2.4	Probabilistic identification	121
6.3	Results	122
6.3.1	Rule-based identification: baseline	122
6.3.2	Rule-based identification: compositionality scores	123
6.3.3	Probabilistic identification	126
6.4	Summary	127
7	Conclusions	130
7.1	Contributions	130
7.2	Future work	133
	References	136
	Index	151
	Appendix	155
A	List of English Compounds	155
B	List of French Compounds	156
C	List of Portuguese Compounds	159

List of Figures

2.1	Representing syntax: constituency tree.	26
2.2	Representing syntax: dependency tree.	26
2.3	Most frequent words in toy corpus.	29
2.4	Visual representation of Pearson's r (left) and Spearman's ρ (right) between human-rated association scores and PMI.	36
2.5	DSM vectors of the target words "fish", "dog" and "cat", considering the contexts "lex" and "swim".	45
2.6	Architecture of word2vec (CBOW and skipgram).	47
3.1	Excerpt from the questionnaire of the compound <i>benign tumor</i> , evaluating compositionality regarding the head of the compound.	56
3.2	Average compositionality (c_H , c_M and c_{WC}) per compound.	58
3.3	Standard deviation (σ_H , σ_M and σ_{WC}) per compound.	59
3.4	Distribution of $c_H \otimes c_M$ according to c_{WC} of each compound.	60
3.5	Compositionality for compounds under different frequency bins.	61
3.6	Compositionality for noun compounds under different PMI bins.	62
4.1	Common compositionality prediction architecture	65
4.2	Number of compositionality prediction models evaluated.	75
5.1	Overall highest results per DSM on English datasets.	78
5.2	Overall highest results per DSM on French and Portuguese datasets.	79
5.3	Best Spearman's ρ per DSM and WINDOW SIZE.	82
5.4	Best Spearman's ρ per DSM and DIMENSION.	83
5.5	Best Spearman's ρ per DSM and WORD FORM	85
5.6	Spearman's ρ for different corpus sizes, running <i>PPMI-thresh</i> (left) and <i>w2v-sg</i> (right).	87
5.7	Distribution of $\text{improv}_{\text{maxsim}}$ as a function of human judgments.	92
5.8	Distribution of $\text{improv}_{\text{maxsim}}$ as a function of <i>uniform</i> scores.	93
5.9	Distribution of $\text{improv}_{\text{geom}}$ as a function of human judgments.	95
5.10	Distribution of $\text{improv}_{\text{geom}}$ as a function of <i>uniform</i> scores.	95
5.11	Compositionality prediction under different frequency bins.	104
5.12	Compositionality prediction under different PMI bins.	105
5.13	Compositionality prediction as a function of human judgments.	107
5.14	Spearman of best <i>uniform</i> models, separated by σ_{WC} ranges.	108
5.15	Compositionality sliding windows, evaluating top w + bottom w compounds, for different values of w .	110
6.1	MWE-annotated output with different match distances.	117

List of Tables

2.1	Most frequent words in toy corpus.	28
2.2	Most frequent word pairs involving “ <i>federal</i> ”.	30
2.3	PMI between “ <i>federal</i> ” and associated nouns.	32
2.4	Human-rated association scores vs PMI.	34
2.5	Human-annotated and system-predicted POS tags.	37
2.6	Binary evaluation measures for POS-tag prediction.	38
2.7	Matrix with targets (rows) \times contexts (columns)	44
3.1	Number of cases of high standard-deviation σ .	58
4.1	228 parameter combinations across all DSMs.	74
5.1	Highest Spearman ρ for all datasets (strict/fallback format for English datasets). The <i>Farahmand</i> dataset uses the highest BF_1 instead.	80
5.2	Results for whole-corpus and parallel predictions.	88
5.3	Spearman scores for best <i>uniform</i> model of each dataset, using different prediction strategies. The <i>Farahmand</i> dataset uses BF_1 .	90
5.4	Highest Spearman score for each prediction strategy individually. <i>Farahmand</i> uses BF_1 .	91
5.5	Average weight of highest-Spearman <i>maxsim</i> model for each dataset.	92
5.6	Outliers regarding positive/negative <i>maxsim</i> improvement.	94
5.7	Outliers regarding positive/negative <i>geom</i> improvement.	96
5.8	Results using a higher number of iterations.	98
5.9	Results for a higher minimum threshold of word count.	99
5.10	Results using a window of size 2+2.	100
5.11	Results for higher numbers of dimensions (<i>PPMI-thresh</i>).	101
5.12	Configurations with highest ρ_{avg} for non-deterministic models.	101
5.13	Intrinsic quality measures for the raw and filtered datasets	102
5.14	Extrinsic quality measures for the raw and filtered datasets	103
5.15	Spearman of best <i>uniform</i> models, separated by c_{wc} ranges.	109
5.16	Comparison of our best models with state-of-the-art results for <i>Reddy</i> .	113
5.17	Comparison of our best models with state-of-the-art BF_1 for <i>Farahmand</i> .	114
6.1	Baseline results for rule-based MWE identifier (English dataset).	122
6.2	Combined baselines for rule-based MWE identifier (English dataset).	123
6.3	Baseline results for rule-based MWE identifier (French dataset).	123
6.4	Results for rule-based MWE identifier using $COMP_{W2V}$.	124
6.5	Results for rule-based MWE identifier using $COMP_{PPMI}$.	124
6.6	Classification of MWE candidates identified by $COMP_{W2V}$.	125

6.7	Results for probability-based MWE identifier (English dataset).	126
6.8	Results for probability-based MWE identifier (French dataset).	127
7.1	Hypotheses evaluated in this thesis.	131

Abbreviations

EN	English language
FR	French language
PT	(Brazilian) Portuguese language
DSM	Distributional Semantic Model
MWE	Multiword Expression
NC	Nominal Compound
NLP	Natural Language Processing
POS	Part of Speech

1. Introduction

The ever-growing presence of computers in our society has put forward an increasing demand for new ways of dealing with human-generated content. The field of Natural Language Processing (NLP) has the goal of mediating this interaction between computers and human language, ranging from the interpretation of written texts on the web to the interaction with spoken commands on hand-held devices. A common theme to many of the NLP tasks is the requirement of *semantic* interpretation (i.e. determining the meaning of the text).

One of the most fundamental assumptions in the field of semantics is that the meaning of a phrase, expression or sentence can be determined from the meanings of its parts. Part of the appeal of this *principle of compositionality*¹ is that it implies that a meaning can be assigned by humans even to sentences that have never been seen before, through the combination of the meaning of familiar words (Goldberg 2015). In the case of NLP, semantic composition can also be an attractive way of deriving the meaning of larger units from their smaller parts. By employing the principle of compositionality, one could design generic NLP systems, able to perform the semantic interpretation of any text.

The representation of the meaning of individual words and their combinations in computational systems has often been addressed by *distributional semantic models* (DSMs). DSMs are based on Harris' distributional hypothesis that the meaning of a word can be inferred from the context in which it occurs (Harris 1954; Firth 1957). In these models, words are usually represented as vectors that, to some extent, capture cooccurrence patterns in corpora. These vectors are assumed to be good proxies for meaning representations. Traditionally, a vector can be built for a target word by explicitly counting all its cooccurrences with context words (Lin 1998; Landauer, Foltz, and Laham 1998). These models, also known as *count-based models* (Baroni, Dinu, and Kruszewski 2014), result in sparse vectors that are often projected into a low-dimensionality space using a statistical technique such as singular value decomposition. The more recent neural-network models, often referred to as *word embeddings*, also represent words as real-valued vectors projected onto some low-dimensional space, but these are obtained as a by-product of training a neural network to learn a function between words and their contexts (Mikolov, Sutskever, Chen, et al. 2013).

Evaluation of DSMs has focused on obtaining accurate semantic representations for single words, and it is on this basis that many optimizations have been proposed (Lin 1999; Erk and Padó 2010; Baroni and Lenci 2010). For instance, state-of-the-art models are already capable of obtaining a high level of agreement with human judgments for predicting synonymy or similarity between single words (Freitag, Blume, Byrnes, et al. 2005; Camacho-Collados, Pilehvar, and

¹ Attributed to Frege (1892/1960).

Navigli 2015; Lapesa and Evert 2017). Although there seems to be a reasonable understanding of the strengths and weaknesses of vector representations for single words, the same is not true for larger units such as sentences. There exist some proposals for modeling the *composition* of individual words to create representations for larger units such as phrases, sentences and even whole documents (Mitchell and Lapata 2010; McCarthy, Keller, and Carroll 2003; Reddy, McCarthy, and Manandhar 2011; Mikolov, Sutskever, Chen, et al. 2013; Ferret 2014). They include the use of simple additive and multiplicative vector operations (Mitchell and Lapata 2010), syntax-based lexical functions (Socher, Huval, Manning, et al. 2012), and the application of matrices and tensors as word-vector modifiers (Baroni and Lenci 2010; Bride, Van de Cruys, and Asher 2015). These operations usually assume the principle of compositionality when building representations for larger units.

However, this assumption is challenged in the case of idiomatic expressions, whose meanings may not be straightforwardly related to their parts (Sag, Baldwin, Bond, et al. 2002). In fact *multiword expressions* (MWEs) display a wide spectrum of idiomaticity, from more compositional to more idiomatic cases (Baldwin and Kim 2010). For instance, although the meaning of *olive oil* can be derived from its parts (as *oil* extracted from *olives*), this is not the case for *snake oil*, which is used to refer to any product of questionable benefit (not necessarily *oil* and certainly not extracted from *snakes*). Such constructions are notoriously challenging for semantically-focused systems, as they are very numerous in a speaker’s lexicon (Jackendoff 1997), but they are often not *compositional*. For example, a non-compositional expression such as *dead end* should not be literally translated into French as **fin morte*, as it would lose its intended meaning. It is therefore crucial to determine to what degree the principle of compositionality applies to a specific expression to ensure its correct semantic interpretation.

The task of *compositionality prediction* consists in assigning a numerical score to a word combination indicating to what extent the meaning of the whole combination can be directly computed from the meanings of its component words. This score can then be used to decide how the combination should be represented in downstream tasks and applications. Given that idiomatic expressions are quite frequent in human languages, compositionality prediction is relevant to any NLP task and application that performs some form of semantic processing. For instance, in machine translation, idiomatic expressions must be translated as an indivisible whole (Cap, Nirmal, Weller, et al. 2015; Salehi, Mathur, Cook, et al. 2015; Carpuat and Diab 2010; Ren, Lü, Cao, et al. 2009). In semantic parsing, one needs to identify complex predicates and their arguments to avoid erroneous analyses (Jagfeld and Plas 2015; Hwang, Bhatia, Bonial, et al. 2010). For word-sense disambiguation, no sense should be ascribed to the individual words pertaining to an idiomatic expression (Schneider, Hovy, Johannsen, et al. 2016; Kulkarni and Finlayson 2011). In all of these cases, there is a need for the preprocessing task of *MWE token identification* in running text, and this opera-

tion may benefit from the availability of compositionality scores.

In this thesis, we discuss approaches for automatically predicting the compositionality of MWEs on the basis of their semantic representation and those of their component words represented using DSMs. To determine to what extent these models are adequate cross-lingually, we evaluate them in three languages, English, French and Portuguese. Since MWEs encompass a large amount of related but distinct phenomena, we focus exclusively on a sub-category of MWEs: *nominal compounds* (NCs).² Nominal compounds (such as *nut case* or *milk tooth*) represent an ideal case study for the work in this thesis, thanks to their relatively homogeneous syntax (as opposed to e.g. verbal idioms such as *take into account*, which may take internal arguments and modifiers), as well as their pervasiveness in the languages under consideration. We assume that, in the future, models able to predict the compositionality of nominal compounds could be generalized to include other categories of MWEs by addressing their morphological and syntactical variability.

By using DSM instances to predict the compositionality of nominal compounds, we are also indirectly evaluating those instances themselves. While evaluations of DSMs based on single words abound, their evaluation on tasks involving MWEs are currently lacking. Some notable exceptions include the works of Reddy, McCarthy, and Manandhar (2011), who compare additive and multiplicative combinations of traditional DSMs, and of Salehi, Cook, and Baldwin (2015) who look at addition-based models for compositionality prediction using both traditional and neural-network DSMs (see Section 4.1 for more details). However, these works do not explore the vast landscape of existing DSM configurations, and may be unable to draw conclusions that are generalizable across languages, DSMs, their parameters and the corpora they are learned on.

The main goal of this thesis is to bridge this gap by presenting a framework for MWE compositionality prediction along with a broad cross-lingual evaluation. We evaluate both intrinsically and extrinsically to what extent DSMs can accurately model the semantics of NCs with various levels of compositionality compared to human judgments. The following section details the contributions that allow us to reach this goal.

1.1. Contributions

The main contributions of this thesis are the following:

Compositionality dataset We construct and evaluate three datasets containing NCs ranging from fully compositional to fully idiomatic. These NCs were manually annotated based on their degree of compositionality. The datasets span multiple languages (English, French and Portuguese), and can be useful in the

² A generalization over compound nouns, see Section 2.3.1.

evaluation of MWE compositionality prediction techniques. Section 3.1 presents a detailed account of the data collection process. The dataset has also been described in a publication (Ramisch, Cordeiro, Zilio, et al. 2016).

Dataset analysis We report the results of a thorough analysis of the three constructed datasets, studying the correlation between compositionality and related linguistic variables. Part of these results has been published in Ramisch, Cordeiro, Zilio, et al. (2016) (focusing on the distribution of annotations) and in Cordeiro, Ramisch, and Villavicencio (2016a) (focusing on inter-annotator agreement). Section 3.2 expands on these results by analyzing the correlation between human-rated scores and distributional characteristics of the NCs.

Compositionality prediction framework We propose a language-independent framework for the prediction of the degree of compositionality in MWE expressions. As part of this framework, we also systematize a set of parameters that can be evaluated across different DSMs, allowing a sound comparison of multiple DSMs under a variety of settings. In Chapter 4, we extend the underlying model with the possibility of six compositionality prediction strategies — two of which are original strategies (*maxsim* and *geom*), proposed in the scope of this thesis. The implementation of this framework is freely available as part of the *mwetoolkit*. The predictive framework has been described in a publication (Cordeiro, Ramisch, and Villavicencio 2016b).

Intrinsic evaluation We evaluate the proposed compositionality prediction model under a variety of settings: different DSMs, different DSM parameters, different corpora parameters, different prediction strategies. Chapter 5 of the thesis extends these results with predictions for Portuguese datasets, including previously unpublished results for one DSM (*lexvec*) and multiple prediction strategies. We additionally evaluate corpus-specific parameters such as corpus size and a new technique of parallel predictions. Furthermore, we consolidate the interpretation of these results through a large set of previously unpublished sanity checks and detailed error analyses. The results have been published in Cordeiro, Ramisch, Idiart, et al. (2016).

Extrinsic evaluation We perform an extrinsic evaluation of the proposed compositionality prediction model by using predicted scores as features in the task of MWE identification. Chapter 6 presents the implementation of an MWE identifier based on syntactic patterns (Cordeiro, Ramisch, and Villavicencio 2016c), and compare its accuracy with the one achieved by a technique of sequence modeling, with and without the help of the predicted scores (Scholivet, Ramisch, and Cordeiro 2017).

1.2. Investigated hypotheses

This thesis investigates a series of hypotheses concerning MWEs and their compositionality, both in relation to the human perception of compositionality and in relation with DSM-based representations. The hypotheses described in this section guide our work — whose main goal is to propose, implement and evaluate distributional methods for the compositionality prediction of MWEs.

Central to this thesis is a framework of compositionality prediction based on DSM representations of semantics. The main assumption behind this framework is that, when the semantics of a compositional MWE can be derived from a combination of its parts, this should be reflected in DSMs. In particular, the vector for the compositional MWE should be similar to the combination of the vectors of its parts. Conversely we can use the lack of similarity between the MWE vector representation and a combination of its parts to detect non-compositionality. We formulate this assumption in the form of the general hypothesis $h_{\text{pred-comp} \approx \text{comp}}$: **MWE compositionality as assessed by human annotators is correlated with compositionality predictions**, where the predictions are based on the distributional representation of MWE elements and MWEs themselves.

We begin our work with the construction of three datasets of nominal compounds with human-annotated compositionality scores. This dataset is analyzed so as to evaluate the hypothesis $h_{\text{idiom} \approx \text{distr}}$: **idiomaticity is correlated with distributional characteristics of MWEs**. In particular, we consider the following sub-hypotheses:

- $h_{\text{idiom} \approx \text{distr.freq}}$ The level of idiomaticity of an MWE is positively correlated with its frequency. The intuition is that exceptional constructions (such as idiomatic MWEs) need to be frequent to ensure their survival in the language (Pinker 1995).
- $h_{\text{idiom} \approx \text{distr.convent}}$ The level of idiomaticity of an MWE is positively correlated with its level of conventionalization. This follows from the literature on MWE type extraction, which uses estimators of conventionalization to identify the idiomatic expressions among a list of MWE candidates (Fazly and Stevenson 2006; Bu, Zhu, and Li 2010; Gurrutxaga and Alegria 2013; El Maarouf and Oakes 2015).

For each dataset, we instantiate DSMs under a variety of configurations, generating a total of more than 8 thousand sets of compositionality predictions. We then evaluate what kinds of variables may influence the accuracy of the highest-ranking configurations. We consider the general hypothesis $h_{\text{accur} \leftarrow \text{MWE}}$: **the accuracy of the model depends on MWE-specific properties**, and we formulate four non-mutually-exclusive sub-hypotheses:

- $h_{\text{accur}} \leftarrow \text{MWE.idiom}$ The accuracy of predicted scores is higher for MWEs that were classified by humans as more compositional (i.e. less idiomatic). The intuition is that DSM representations should be more faithful to the reality for compositional MWEs, which follow the regularities that are normally exploited in other works in the literature (Mitchell and Lapata 2010; Mikolov, Sutskever, Chen, et al. 2013).
- $h_{\text{accur}} \leftarrow \text{MWE.diffic}$ The accuracy of predicted scores is lower for MWEs that are more difficult to annotate for humans, as measured through the level of agreement among annotators. The intuition is that one would expect the predictive model to have difficulty in the same MWEs that posed a problem for humans, either due to some inherent difficulty in the MWE or due to less reliability of the data.
- $h_{\text{accur}} \leftarrow \text{MWE.freq}$ The accuracy of predicted scores is positively correlated with the frequency of the MWE in the corpus. The intuition is that low-frequency expressions should have a less trustworthy representation inside DSMs.
- $h_{\text{accur}} \leftarrow \text{MWE.convent}$ The accuracy of predicted scores is positively correlated with the conventionalization of the MWE. The intuition is that the elements of highly conventionalized MWEs are more likely to share contexts, and thus have a more compatible vector representation.

In addition to these MWE-centric hypotheses, we consider the specific choice of DSMs and internal parameters in the task of compositionality prediction. The hypothesis $h_{\text{accur}} \leftarrow \text{DSM}$ is that the **accuracy of the model depends on DSM-specific parameters**. In particular, we consider two sub-hypotheses:

- $h_{\text{accur}} \leftarrow \text{DSM.window}$ The accuracy of compositionality prediction depends on the amount of content that is taken into account at each occurrence of a word (i.e. the size of the context window). More context should lead to a more precise representation and thus result in better predictions.
- $h_{\text{accur}} \leftarrow \text{DSM.dims}$ The accuracy of compositionality prediction depends on the number of dimensions in each vector generated by the DSM. A higher number of dimensions should allow for a more fine-grained representation of the data.

Along with an influence from DSM-specific parameters, we also consider the impact of different corpus-specific configurations. The hypothesis we evaluate is $h_{\text{accur}} \leftarrow \text{corpus}$: **accuracy of the model depends on corpus-specific parameters**. We evaluate the following sub-hypotheses:

- $h_{\text{accur}} \leftarrow \text{corpus.wordform}$ Higher-quality compositionality predictions are obtained when the corpus is preprocessed so as to reduce the sparseness of the word occurrences (e.g. through lemmatization). The intuition is that the reduction in sparseness should allow DSMs to generate vectors from a more varied number of contexts, and that these vectors would thus be more robust than the ones generated without preprocessing.
- $h_{\text{accur}} \leftarrow \text{corpus.size}$ DSM representations built from larger corpora outperform representations built from smaller corpora. The intuition is that more varied occurrences of each word allow the construction of DSM representations that more faithfully correspond to the actual semantics.
- $h_{\text{accur}} \leftarrow \text{corpus.parallel}$ Multiple parallel DSM representations built from different parts of the corpora can be combined to achieve equivalent compositionality predictions. The intuition is that these representations could still provide a high-quality description of the underlying semantics, while allowing for the final predictions to be calculated through a combination of multiple computational resources (e.g. computer clusters).

Regarding the compositionality prediction model, we consider six different predictive strategies (defined in Section 5.4). The hypothesis h_{strat} is that the **accuracy of the model depends on the predictive strategy**. In particular, we consider these three sub-hypotheses:

- $h_{\text{strat.partial-info}}$ Predictions derived only from parts of an MWE (i.e. its syntactic head) will be less accurate than predictions that consider all words in the MWE. The intuition is that, by using only part of the available distributional information, these strategies are limited in their ability of predicting the compositionality of the MWE as a whole.
- $h_{\text{strat.maxsim}}$ We can improve the score prediction of compositional MWEs through a strategy that favors compositional interpretations. More specifically, we can improve predictions if we assign weights to the vector representation of each member word of an MWE so as to maximize its compositionality score. The *maxsim* strategy proposed in this thesis models this assignment of weights favoring composition.
- $h_{\text{strat.geom}}$ We can improve the score prediction of idiomatic MWEs through a strategy that favors idiomatic interpretations. More specifically, we propose the *geom* strategy that multiplies individual predictions of compositionality for all words in an MWE. If any word has been identified as idiomatic (i.e. having a small level of compositionality), the compositionality score of the MWE as a whole will be reduced.

One final hypothesis we consider is related to the use of predicted compositionality scores in an extrinsic evaluation. Here, we consider the task of MWE token identification, which can be seen as an important preprocessing step to semantic applications (such as machine translation or text simplification). The relevant hypothesis is $h_{\text{pred-comp}} \rightarrow \text{ident-accur}$: **predicted compositionality scores are useful in the task of MWE identification**. This hypothesis is evaluated by a comparison of the accuracy of the MWE identification with and without the use of predicted compositionality scores.

1.3. Context of the thesis

The work in this thesis has been produced as part of a joint supervision (co-tutelle) between two universities: Universidade Federal do Rio Grande do Sul (UFRGS, Brazil) and the Aix-Marseille Université (AMU, France). I have spent part of my research time in each university, and most of the publications derive from this cooperation.

This thesis is centered on the topic of multiword expressions (MWEs) and distributional semantic models (DSMs), which are two areas of research that have been growing in the latest decade. The increasing interest in MWE research has been the main motivation behind the formation of the PARSEME network of researchers (ICT COST Action IC1207). Some of the work done in conjunction with PARSEME researchers, notably the shared task on verbal MWE identification (Savary, Ramisch, Cordeiro, Sangati, et al. 2017), falls out of the scope of this thesis, but has been paramount to a broadening of my view of the field. Moreover, part of the work described in this thesis was carried out in the context of PARSEME-FR (ANR-14-CERA-0001), a French-language spin-off of PARSEME.

This thesis is also closely related to a software called *mwetoolkit* (Ramisch 2015), which was both used and extended during this thesis. The *mwetoolkit* is one of the major resources published by the research group in the side of UFRGS, and it is currently being maintained in joint work with Carlos Ramisch in AMU. The contributions from this thesis are integrated and freely available as part of the *mwetoolkit*.³

1.4. Publications

I present below a list of papers that have been published or accepted for publication during the period of my PhD studies. The papers that are directly relevant to the topic of this thesis are the following ones:

- Cordeiro, Ramisch, Idiart, Villavicencio. *Predicting the Compositionality of Nominal Compounds: Giving Word Embeddings a Hard Time*. In: ACL 2016.

³ <http://mwetoolkit.sf.net>

- Ramisch, Cordeiro, Zilio, Idiart, Villavicencio, Wilkens. *How Naked is the Naked Truth? A Multilingual Lexicon of Nominal Compound Compositionality*. In: ACL 2016.
- Cordeiro, Ramisch, Villavicencio. *mwetoolkit+sem: Integrating Word Embeddings in the mwetoolkit for Semantic MWE Processing*. In: LREC 2016.
- Cordeiro, Ramisch, Villavicencio. *Filtering and Measuring the Intrinsic Quality of Human Compositionality Judgments*. In: MWE 2016.
- Cordeiro, Ramisch, Villavicencio. *UFRGS&LIF: Rule-Based MWE Identification and Predominant-Supersense Tagging*. In: SemEval 2016.
- Scholivet, Ramisch, Cordeiro. *Sequence Models and Lexical Resources for MWE Identification in French*. In: PMWE 2017. Submitted for review.
- Wilkens, Zilio, Cordeiro, Paula, Ramisch, Idiart, Villavicencio. *LexSubNC: A Dataset of Lexical Substitution for Nominal Compounds*. In: IWCS 2017.
- Cordeiro, Ramisch, Villavicencio. *Token-based MWE Identification Strategies in the mwetoolkit*. In: PARSEME 2015.
- Cordeiro, Ramisch, Villavicencio. *Nominal Compound Compositionality: A Multilingual Lexicon and Predictive Model*. In: PARSEME 2016.
- Cordeiro, Ramisch, Villavicencio. *MWE-aware corpus processing with the mwetoolkit and word embeddings*. In: W-PROPOR 2016.

Other papers published (or submitted for review) during this period are not directly implicated in the present thesis. Nevertheless, they all have in common the topic of MWE processing and semantic representation:

- Savary, Ramisch, Cordeiro, Sangati, Vincze, QasemiZadeh, Candito, Cap, Giouli, Stoyanova, Stoyanova, Doucet. *The PARSEME Shared Task on Automatic Identification of Verbal Multiword Expressions*. In: MWE 2017.
- Zilio, Wilkens, Möllmann, Wehrli, Cordeiro, Villavicencio. *Joining forces for multiword expression identification*. In: PROPOR 2016.
- Savary, Ramisch, Cordeiro, Candito, Vincze, Sangati, QuasemiZadeh, Giouli, Cap, Stoyanova, Stoyanova, Doucet. *The PARSEME shared task on automatic identification of verbal multiword expressions*. In: PMWE 2017. Submitted for review.
- Savary, Candito, Mititelu, Bejček, Cap, Čéplö, Cordeiro, Eryiğit, Giouli, Gompel, Hachohen-Kerner, Kovalevskaite, Krek, Liebeskind, Monti, Escartín, Plas, Qasemizadeh, Ramisch, Sangati, Stoyanova, Vincze. *PARSEME multi-lingual corpus of verbal multiword expressions*. In: PMWE 2017. Submitted for review.

1.5. Thesis structure

The remainder of this thesis is structured as follows:

- Chapter 2 presents the terminology and necessary background in the statistical representation of languages. It also presents a literature review on MWE research, as well as related work on the semantic representation of words and MWEs.
- Chapter 3 describes the methodology used in the construction of three new datasets of nominal compounds annotated with compositionality scores. It also presents an analysis of the score distribution and difficulty of annotation, as well as the correlation between the annotated score and distributional characteristics of the MWEs.
- Chapter 4 presents our framework of MWE compositionality prediction, with a detailed description of the experimental setup that will be used in the following Chapter.
- Chapter 5 presents a large-scale intrinsic evaluation of our model of compositionality prediction against datasets of human-rated compositionality scores, focusing on nominal compounds as a specific category of MWE.
- Chapter 6 presents an extrinsic evaluation of our model of compositionality prediction, in the form of its application in the task of MWE identification.
- Chapter 7 presents some conclusions and perspectives of future work.

2. Background

This chapter describes the background information that is essential for the remainder of the thesis. Section 2.1 defines the basic terminology that will be used in the rest of the thesis. Section 2.2 presents the use of statistics in NLP, going from co-occurrence measures to evaluation methods. Section 2.3 then presents the motivation and challenges associated with the research on MWEs. Finally, Sections 2.4 and 2.5 present an overview of lexical semantic representations in NLP, on the level of single words and MWEs. Since it constitutes the core of this thesis, the state of the art on compositionality prediction will be presented later, in Section 4.1, along with the definition of the model that we propose and evaluate.

2.1. Basic terminology

The work in this thesis focuses on the interpretation of written texts. Linguists often classify the interpretation of texts into different layers of abstraction. For example, consider the sentence *“this student published a paper”*. These are some of the levels in which this sentence can be analyzed:

- **Morphology:** A word can change its base form to express a change of grammatical category (derivational morphology; e.g. *“publish”* → *“publisher”*) or to represent a variation such as gender, number or tense (inflectional morphology; e.g. *“publish”* → *“published”*).
- **Syntax:** Just like morphology works on the level of words and their internal structure, syntax can be used to analyze sentences based on their internal structure (word order). Such analysis, when applied to the above sentence, could inform us that *“published”* is the main verb, and that *“this student”* and *“a paper”* are respectively the subject and direct object of this verb.
- **Semantics:** The field of semantics looks at the meaning of words and phrases. For example, although the word *“paper”* is often used as reference to the material that is created from cellulose, in this particular context it should be interpreted as a scientific document (e.g. it can be in electronic format).
- **Pragmatics:** While one may be able to build an accurate abstract mental model for a word such as *“student”* given only the surrounding words, the specific reference in *“this student”* can only be resolved by taking the underlying context into account. If this sentence is uttered by a person who is pointing at someone else, for example, one may conclude through pragmatics who is the referent of *“this student”*.

In the area of NLP, the operations that can be performed on texts are usually conceptually organized in tasks. Many of these tasks may take as an input a collection of sentences and have as an output the same sentences with an added layer of information (Jurafsky and Martin 2008). Tasks are often performed in succession as a pipeline, where the output of a task is the input for the next one. These are some of the tasks that are often performed in any end-to-end application in NLP involving text analysis:

- **Tokenization:** The goal of tokenization is to break a written text into *tokens*, which correspond somewhat to the linguistic notion of *words*. While the precise definition of word depends on language-specific semantics, a *token* is a pragmatic concept used in NLP, and its definition will often vary depending on the tokenizer at hand. For example, the sentence

Mr. Smith doesn't eat bananas with a fork.

could be tokenized as such:

Mr. Smith does n't eat bananas with a fork .

The punctuation associated with the abbreviations (e.g. “*Mr.*”) is often tokenized along with the preceding characters, while the period at the end of the sentence is considered a token apart. Also note that the contraction “*doesn't*” may be separated in two tokens, indicating the two underlying words “*do+not*”. The use of tokens allows a trade-off between linguistic accuracy and practicality of implementation.

- **POS tagging:** In a given context, every word can be associated with a grammatical class, known as its Part-of-Speech (POS). The task of POS tagging consists in identifying the POS of each token in the text according to a *tagset*. A list of POS tags for the sentence above could be the following one:

PROPN PROPN AUX PART VERB NOUN ADP DET NOUN PUNCT

These tags follow the Universal POS tagset, which standardizes a common set of tags across a variety of languages (Petrov, Das, and McDonald 2011).

- **Lemmatization:** The goal of lemmatization is to identify a canonical form for the tokens in the text. For each token in its *surface form*, the lemmatizer will produce the corresponding *lemma*. For the example above, these could be the resulting lemmas:

Mr. Smith do not eat banana with a fork .

Lemmatization is usually applied to neutralize distinctions caused by morphological inflection (e.g. “*publish/publishes/published/publishing*” → “*publish*”).

- **Parsing:** In the goal of understanding a sentence, it is often useful to have a representation of how words are grouped to form larger structures. This syntactic information can be encoded in different ways, and it usually corresponds to one of two classes of grammar theories: phrase structure and dependency grammar.

Figure 2.1 presents a syntax tree for the sentence “*this man eats bananas with a fork*” according to the phrase structure grammar. In this representation, words that behave syntactically as a single unit are grouped as *phrases* (also known as constituents). For example, the phrase “*this man*” behaves as a single noun, and is thus grouped under a single node of the tree, known as a noun phrase (NP).

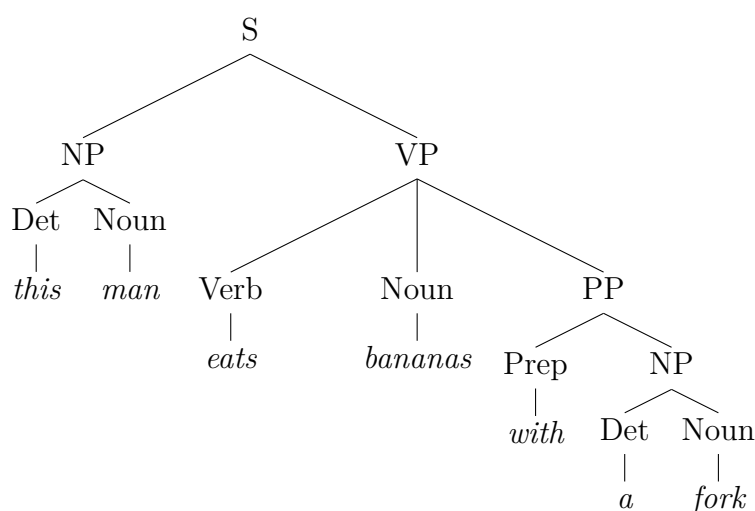


Figure 2.1.: Representing syntax: constituency tree.

In a dependency tree representation, words are connected by labeled edges, in what is known as a dependency relation (see Figure 2.2). This relation connects *heads* (e.g. a verb) to their *dependents* (e.g. a verb’s subject). The root of the tree is usually the main verb (in the example, the verb *eats*), and everything else is connected to this verb through a chain of dependency relations.




Figure 2.2.: Representing syntax: dependency tree.

Many works in NLP are centered around the idea of a *corpus*. A corpus is a large body of (written or spoken) text that can be considered representative of a human language (Mitkov 2005). In NLP, this is often taken to mean that corpora should be as large as possible. The construction of a corpus is often performed automatically (e.g. by crawling the web), digitally storing the data according to some file format for further processing. In the case of written corpora, the pre-processing pipeline may include tasks such as the aforementioned tokenization, POS tagging and lemmatization.

One important distinction that must be made in NLP is the one between *types* and *tokens*. While a given corpus may have many concrete instances of a given word (tokens), it is often useful to talk about the unique concept that instantiated those tokens (the type). In the same manner that tokens can be assembled in sentences to form corpora, types can be assembled in dictionary entries. Such dictionaries are also known as lexicons, or more generally as lexical resources. The set of all types instantiated in a corpus is known as its *vocabulary*, which we will denote in this thesis as V .

Finally, most applications in NLP involve some kind of prediction, where the computational system attempts to replicate the outcome that a human would have produced. This prediction is then compared against a blind *test set*. Test sets are often handcrafted based on direct human knowledge, which forms an authoritative sample of test cases and their solutions (in which case it is also known as a *gold standard*). Some systems are designed so as to perform their predictions solely based on a programmed algorithm. Others, known as *supervised* approaches, require a *training set* with examples from which the system may learn to perform the predictions (Manning and Schütze 1999; Jurafsky and Martin 2008)

2.2. Language and statistics

Now that we have defined the basic terminology, we turn to a statistical view of the properties of human language. We start with some background on the occurrence counts of isolated words (Section 2.2.1) and of word pairs in the same context (Section 2.2.2). We then present measures of association between word pairs (Section 2.2.3). We also consider the creation of human-rated datasets, focusing on measures of inter-rater agreement (Section 2.2.4). Finally, we consider different ways of measuring the accuracy of computational predictions of human annotation, both in the case of continuous data (Section 2.2.5) and in the case of categorical data (Section 2.2.6).

2.2.1. Occurrence counts

Many properties of words can be inferred from their statistical behavior with regards to human language. One of the simplest statistical measures that can be considered is the number of *occurrences* of each word in a corpus. For example, consider the following three toy corpora: English-language (EN) Wikipedia entry “U.S.A”, French-language (FR) Wikipedia entry “France” and Portuguese-language (PT) Wikipedia entry “Brasil” (*Brazil*).¹ Table 2.1 presents the tokens with highest number of occurrences (#occur) in each of these toy corpora.

Rank	#occur	EN word	#occur	FR word	#occur	PT word
1	1256	the	1340	de (<i>of</i>)	770	de (<i>of</i>)
2	742	of	912	la (<i>the</i>)	460	e (<i>and</i>)
3	701	and	799	et (<i>and</i>)	401	do (<i>of+the</i>)
4	473	in	533	le (<i>the</i>)	349	a (<i>the/to</i>)
5	268	The	480	des (<i>the/of+the</i>)	323	o (<i>the</i>)
6	245	United	423	en (<i>in</i>)	233	em (<i>in</i>)
7	235	to	392	les (<i>the</i>)	204	da (<i>of+the</i>)
8	206	a	379	du (<i>of the</i>)	151	Brasil (<i>Brazil</i>)
9	191	States	365	à (<i>to</i>)	144	no (<i>in+the</i>)
10	186	is	286	est (<i>is</i>)	136	que (<i>that</i>)

Table 2.1.: Most frequent words in toy corpus.

Two notable categories of words can be seen in the table:

- Words that have been biased by the input corpus (e.g. *States*). These words are an artifact of the chosen corpora, and would have a notably lower ranking in a more general corpus.
- Words that only convey a general meaning (e.g. prepositions, articles), serving mainly to connect other words according to the underlying grammar. These words are known as function words, and contrast with the more semantically distinguished content words (e.g. nouns, verbs). In NLP, a similar concept to function words is the one of *stopwords*. When dealing with semantic tasks, one very common step of corpus preprocessing involves the removal of stopwords from the text. The definition of stopwords may be based on their POS tag (e.g. removing all prepositions from the text) or based on a list of e.g. the 50 or 100 words with the highest number of occurrences in the corpus.

When looking at the number of occurrences per word, we can also see a clear pattern: the number of occurrences of each word decreases proportionally to

¹ Wikipedia entries collected on 2017-08-10.

its ranking. This pattern, known as Zipf's Law, can more easily be seen in a graph, as in Figure 2.3, which shows the number of occurrences of the top 50 most frequent words in the three Wikipedia pages. If we consider the number of occurrences for the English word at rank 5, we see that it is about 5 times smaller than the number of occurrences for the word at rank 1. Similarly, the word at rank 10 appears about 10 times less often in the corpus.




Figure 2.3.: Most frequent words in toy corpus.

The above effects can be observed in all human languages, even in these severely small corpora (e.g. the English toy corpus has less than 20 thousand tokens). When comparing across different languages, one might be tempted to consider the number of occurrences itself, but this number presents a pitfall: it is directly dependent on the size of the corpus. An alternative solution is the use of *frequency*, defined as the ratio between the number of occurrences and the size of the underlying corpus:

$$\text{freq}(w) = \frac{\#\text{occur}(w)}{\text{corpus size}}.$$

Frequency values range between 0.0 and 1.0 regardless of corpus size, and can be thus regarded as a normalized version of the number of occurrences.

2.2.2. Co-occurrence counts

Similarly to how one may derive statistical information from the occurrences of isolated words, one may consider the statistical properties of pairs of words that *co-occur* (i.e. that occur in the same context). Co-occurrence may be calculated based on syntactic information (i.e. two words co-occur if they are connected in

a dependency tree), or on a sliding window of adjacent words (i.e. two words are considered to co-occur if there are at most k other words in between). For example, Table 2.2 presents the top 10 most frequent co-occurrences of the pattern “*federal* <noun>” in the toy corpus, along with their co-occurrence counts (#co-occur).² In this thesis, we will refer to every word of interest (e.g. “*federal*”) as a *target word*. The words that co-occur with a target word will be called *context words* (e.g. “*government*”, “*courts*”, etc).

This type of word-pair information is often structured as a *co-occurrence matrix*. For a vocabulary V of size $|V|$, a co-occurrence matrix of dimension $|V| \times |V|$ contains the number of occurrences of each (target, context) pair as seen in the corpus. One can think of these matrices as big square tables of numbers, with $M_{i,j}$ representing the value at row i and column j . Words are mapped to row/column indexes (e.g. “*government*” could be word_{53}), so that both row_i and column_i refer to word_i . The co-occurrence between word_j and word_k can then be obtained at the matrix position $M_{j,k}$. Due to the symmetry of this statistical relation, the matrix position $M_{k,j}$ yields the same result.

#co-occur	word pair
9	federal government
4	federal courts
4	federal district
4	federal level
4	federal taxes
2	federal debt
2	federal law
2	federal outlays
2	federal republic

Table 2.2.: Most frequent word pairs involving “*federal*”.

According to the distributional hypothesis, the meaning of words is associated with the context that they share. For the example in Table 2.2, all of these nouns have the word “*federal*” as a context in common, from which we can infer that they have related meaning. Moreover, the co-occurrence of two words can be seen as an evidence that they are semantically “close” in a sense — e.g. these nouns are closer in meaning to the word “*federal*” than other nouns not shown here. The underlying idea is that semantically related words tend to be used together, and hence words that often appear together can be assumed to be semantically related. For example, when talking about a “*government*”, one will often use the word “*federal*”. Similarly, among all things “*federal*”, the “*government*” is a notable example that will come to a speaker’s mind. Co-occurrence

² Extracted using the sliding window method with $k = 0$ (i.e. only adjacent word pairs were considered)

counts can also be used to build vector models of word semantics (as presented in Section 2.4.2).

2.2.3. Association measures

Lexical co-occurrence can be considered a rudimentary way of estimating the level of *association* between a pair of words. Word association is a statistical property that can indicate the predictability of a specific combination of words. For example, the word “*population*” appears 41 times in the text, but since none of these is paired up with the word “*federal*”, it is not taken into account when representing its semantics through co-occurrence, suggested that these two words are not associated.

There is a downside to the use of co-occurrence counts as measures of association: some word pairs may co-occur with high frequency solely due to the fact that one of the words in the pair is frequent itself. For example, even in a syntactically restricted word-pair pattern such as “*federal* <word>”, the toy corpus contains 2 occurrences of “*federal and*”, in the coordinations “*federal and state*” and “*federal and military*”. The word “*and*” is not particularly meaningful to the definition of “*federal*”, and its co-occurrence with “*federal*” is due to its high overall frequency (788 occurrences in the toy corpus).

Other measures can better capture the association between word-pairs by reducing the effect of words that have high isolated corpus occurrence counts. One such measure is the pointwise mutual information (PMI), which can be defined as

$$\text{PMI}(w_1, w_2) = \log \left(\frac{\text{freq}(w_1 w_2)}{\text{freq}(w_1) \cdot \text{freq}(w_2)} \right).$$

PMI considers the co-occurrence frequency as well as the individual frequencies of the words. Its value can range from negative infinity to positive infinity. In an extreme case, if the word w_1 always appears alongside the word w_2 in the corpus, $\text{freq}(w_1) = \text{freq}(w_1 w_2)$, and thus $\text{PMI} = \log \left(\frac{1}{\text{freq}(w_2)} \right) = \log \left(\frac{\text{corpus size}}{\#\text{occur}(w_2)} \right)$, which can reach high values for moderate numbers of occurrence of w_2 . On the other hand, if w_1 appears too often by itself and not as often in conjunction with w_2 , the denominator will be considerably high, and thus the PMI between the two words will be the logarithm of a low number — i.e. a low number itself (Church and Hanks 1990).

Table 2.3 presents the number of individual occurrences, word-pair occurrence and PMI for the word pairs mentioned in Table 2.2. Note how the PMI is close to zero for the non-associated word pair “*federal and*”, but is much higher when calculated between the word “*federal*” and highly associated nouns (e.g. “*government*”). Moreover, even among these associated pairs, we can see a distinction in the strength of PMI: the highest score belongs to “*federal outlays*”, reflecting the fact that “*outlays*” only appears in the corpus in conjunction with the word

“*federal*”. A significantly lower score is obtained by “*federal law*”, as the word “*law*” often appears in other contexts inside the corpus (with only 2 out of the 19 occurrences being the expression “*federal law*”).

word ₁	word ₂	#occur(w_1)	#occur(w_2)	#co-occur	PMI
federal	government	54	38	9	4.4
federal	courts	54	7	4	5.3
federal	district	54	10	4	4.9
federal	level	54	16	4	4.4
federal	taxes	54	15	4	4.5
federal	debt	54	9	2	4.3
federal	law	54	19	2	3.6
federal	outlays	54	2	2	5.8
federal	republic	54	5	2	4.9
federal	and	54	788	2	-0.1

Table 2.3.: PMI between “*federal*” and associated nouns.

The frequency value $\text{freq}(w_k)$ can be interpreted as the probability of a random word being precisely w_k . As a consequence of this, negative PMI values imply that the word-pair is occurring less often than would be predicted by chance (through the multiplication $\text{freq}(w_1) \cdot \text{freq}(w_2)$, which estimates the expected probability of two independent words co-occurring). However, such negative scores are often deemed unreliable (Jurafsky and Martin 2008). A common solution is to use Positive PMI (PPMI), which eliminates the negative scores:

$$\text{PPMI}(w_1, w_2) = \max\{0, \text{PMI}(w_1, w_2)\}.$$

Association measures are often used for the estimation of the level of *conventionalization* of expressions. Conventionalization refers to the degree to which the specific choice of words and word order can be seen as fixed in the language, when referring to a particular concept. For example, while one may talk about the “*forecast of the weather*”, there is a distinctive markedness to this choice of words, and native speakers are more likely to refer to this concept as “*weather forecast*” instead. Constructions such as the latter are thus said to be more conventionalized. They have also been described as *collocations*, i.e. statistically idiosyncratic word combinations (Farahmand, Smith, and Nivre 2015; Baldwin and Kim 2010).

Other association scores include the Dice coefficient, Student’s t -test, χ^2 tests, the log-likelihood ratio test (Dunning 1993) and NPMI (Bouma 2009). A thorough description of commonly used association measures, along with their evaluation in the context of collocation extraction, can be found in Pecina (2010).

2.2.4. Inter-rater agreement measures

One of the goals of NLP is to be able to replicate human-like processing of natural language. To this end, datasets containing human annotation of some linguistic phenomenon are often useful in the tuning and evaluation of NLP systems. The construction of these datasets may be performed by experts (with appropriate linguistic background), or may be framed in such a way that the task may be performed by laypeople. In both cases, there must be a set of guidelines, which instruct the annotator on what must be annotated and provide a solution to common corner cases.

Human language can be often ambiguous (especially when it comes to semantics), and this ambiguity may cause disagreement of annotation even among highly well-trained experts. One indicator of the quality and objectivity of both the annotation guidelines and the resulting dataset could be the fraction of annotations in which all annotators agree:

$$p_{\text{agree}} = \frac{\text{cases of agreement}}{\text{all annotations}}$$

In the case of two annotators, another solution would be to calculate the linear correlation and categorical evaluation measures (see Section 2.2.5).

However, these measures do not take into account the probability of *chance agreement* among annotators. For example, consider a POS-tag annotation task with an inventory of 17 possible POS tags. Even if two annotators decide to assign POS tags randomly to every word in the corpus, they will still be in agreement in around $p_{\text{chance}} = \frac{1}{17} = 0.058 = 5.8\%$ of the cases. An alternative measure that does take this probability of chance agreement into account is Cohen's *kappa* coefficient (Artstein and Poesio 2008):

$$\kappa = \frac{p_{\text{agree}} - p_{\text{chance}}}{1 - p_{\text{chance}}}$$

Kappa scores are usually lower-bounded by 0.0 (which indicates pure chance agreement), and are always upper-bounded by 1.0 (which indicates perfect agreement). The kappa score can be generalized for more than two annotators by using an appropriate estimate of p_{chance} .

One downside of the kappa coefficient is the fact that it is restricted to categorical data, with the same weight applied to all disagreements. An alternative measure that calculates multi-annotator agreement while taking into account the distance between ordinal ratings is Krippendorff's *alpha* score (Artstein and Poesio 2008):

$$\alpha = 1 - \frac{D_{\text{disagreement}}}{D_{\text{chance disagreement}}},$$

in which the measure of disagreement D is an average of the variance in the

distance between ordinal ratings for all paired-up annotators.

2.2.5. Evaluation measures for continuous data

When dealing with large amounts of data, it is useful to have a way of summarizing results. In particular, one often needs to measure how similar are two sets of numerical scores. For example, consider Table 2.4, in which each word pair has been manually annotated by a human rater based on the perceived level of association between the words. While the human and PMI scores of each word pair can be individually compared with ease, it is not immediately obvious whether the PMI values are a good estimation of the human ratings. The similarity of these two variables can be calculated through *correlation* measures.

word ₁	word ₂	Human score	PMI score
federal	government	6.0	4.4
federal	courts	8.8	5.3
federal	district	3.0	4.9
federal	level	0.7	4.4
federal	taxes	5.0	4.5
federal	debt	2.2	4.3
federal	law	1.5	3.6
federal	outlays	9.5	5.8
federal	republic	8.0	4.9
federal	and	0.0	-0.1

Table 2.4.: Human-rated association scores vs PMI.

One way of calculating the correlation between two paired datasets X and Y (each with N elements) is through the covariance:

$$\text{cov}(X, Y) = \frac{\sum_{i=1}^N (X_i - \bar{X})(Y_i - \bar{Y})}{N - 1}$$

The covariance between human and PMI scores in Table 2.4 is 3.86. Covariance values close to zero serve as an evidence that the datasets are not correlated: a positive deviation from the mean in the first dataset ($X_i - \bar{X}$) is equally likely to be paired up with a positive and a negative deviation in the second dataset ($Y_i - \bar{Y}$). Covariance values distant from zero (whether negative or positive) indicate a tendency of both datasets towards higher deviations for the same data points. The covariance score is upper-bounded by the max of the variances $\sigma^2(X)$ and $\sigma^2(Y)$, and is similarly lower-bounded by the min of $-\sigma^2(X)$ and $-\sigma^2(Y)$. The variance itself is calculated as a special case of the covariance:

$$\sigma^2(X) = \text{cov}(X, X).$$

For the values in Table 2.4, the variance in human scores was 12.16, while the variance in PMI was 2.64. Variance scores are either 0.0 or positive, with higher values representing data that differs more heavily from the average.

Covariance scores are generally deemed hard to interpret, as they depend on how the datasets themselves deviate from the mean. One way of dealing with this difficulty is to normalize the covariance based on the variance of both datasets. This is what is done in *Pearson's r* coefficient:

$$r(X, Y) = \frac{\text{cov}(X, Y)}{\sqrt{\sigma^2(X)} \cdot \sqrt{\sigma^2(Y)}},$$

This normalization permits a more straightforward interpretation of the correlation, as Pearson scores range from -1.0 (perfect negative correlation) to 1.0 (perfect positive correlation). A neutral score of 0.0 represents no correlation, and values in between represent partial levels of correlation. The Pearson coefficient between human and PMI scores in Table 2.4 is 0.68 , indicating that the values are highly correlated.

Pearson scores measure the linear correlation between paired-up values from two datasets. Sometimes, it may be considered more appropriate to interpret these values as an indicator of the *ranking* of two data items instead. In this case, the correlation between the values themselves may not be as interesting as the ability of both sets of scores to rank the items in the same order. Thus, an alternative measure of correlation is *Spearman's ρ* rank coefficient, which represents the linear correlation between the *ranks* from two datasets. Spearman's ρ can be defined in terms of Pearson's r :

$$\rho = r(\text{rank}(X), \text{rank}(Y)),$$

where the rank operation maps increasing values in the dataset to consecutive integer ranks. For example, for the human scores in Table 2.4, $\text{rank}(X)$ would map $[0.0, 0.7, 1.5 \dots 9.5] \mapsto [\#1, \#2, \#3 \dots \#10]$. Similarly, $\text{rank}(Y)$ would map PMI scores as $[-0.1, 3.6, 4.6 \dots 5.8] \mapsto [\#1, \#2, \#3 \dots \#10]$. As in the case of Pearson, Spearman scores range from -1.0 (perfect negative rank correlation) to 1.0 (perfect positive rank correlation).

Figure 2.4 presents a graphical visualization of these two approaches:

- On the left, we see the PMI values as a function of human ratings themselves. We can visually confirm a tendency of PMI scores to be higher for higher values of human rating. Pearson's r calculates the level of this correlation between the values themselves (visually: how much the points in this graph resemble a straight line). The Pearson coefficient in this example is $r = 0.68$.
- On the right, we see the PMI score ranks as a function of the rank of the human ratings. Values in the x-axis were ordered based on successive ranks

(e.g. $0.0 \mapsto \#1$, $0.7 \mapsto \#2$, and so on). Successive ranks in the y-axis are indicated in little squares (e.g. $5.3 \mapsto \#9$).³ In both axes, we see that consecutive points are separated by an equal amount of space, regardless of the actual value, as each point is plotted based on the successive integer ranks. Spearman’s ρ calculates the linear correlation between these ranks (visually: how much the points in this graph resemble a straight line). The Spearman rank coefficient in this example is $\rho = 0.87$.




Figure 2.4.: Visual representation of Pearson’s r (left) and Spearman’s ρ (right) between human-rated association scores and PMI.

2.2.6. Evaluation measures for categorical data

Correlation measures are able to adequately compare two sets of items with quantitative (e.g. real-valued) scores. However, they are unfit for comparisons involving qualitative categorical scores. The need for comparing categorical scores arises, for example, in the task of POS tagging (see Section 2.1). Consider this sample excerpt from the English toy corpus: “*The first inhabitants of North America migrated from Siberia by way of the Bering land bridge*”. Table 2.5 presents an example of automatic system prediction of POS tags along with the corresponding human-annotated POS tags.

Consider the human and system annotations of the NOUN tag. Humans annotated 4 occurrences (“*inhabitants*”, “*way*”, “*land*”, “*bridge*”), while the system predicted 3 occurrences (“*inhabitants*”, “*North*”, “*bridge*”). Out of these, 2 occurrences were *true positives*, i.e. system predictions that matched human annotations (“*inhabitants*” and “*bridge*”). One way of summarizing the predictive

³ Tie-breaking is not performed in this example for 4.4 and 4.9. In the remainder of the thesis, tie-breaking is performed by assigning the average of all tied ranks (e.g. $4.9 \mapsto \#7.5$).

Word	Human POS	System POS
The	DET	DET
first	ADJ	ADJ
inhabitants	NOUN	NOUN
of	ADP	ADP
North	PROPN	NOUN
America	PROPN	PROPN
migrated	VERB	VERB
from	ADP	ADP
Siberia	PROPN	PROPN
by	ADP	ADP
way	NOUN	ADV
of	ADP	ADP
the	DET	DET
Bering	PROPN	PROPN
land	NOUN	VERB
bridge	NOUN	NOUN

Table 2.5.: Human-annotated and system-predicted POS tags.

quality of this system for NOUN tags is by means of *precision* and *recall*:

$$\text{precision} = \frac{\text{true positives}}{\text{total system predictions}}, \quad \text{recall} = \frac{\text{true positives}}{\text{total human annotations}}.$$

The precision measures how many of the predictions were correct, while the recall measures how many of the tags were correctly predicted. These two measures are often further combined into a single score using the harmonic mean, yielding the F_1 score:

$$F_1 = \frac{2}{\frac{1}{\text{precision}} + \frac{1}{\text{recall}}} = \frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}.$$

Table 2.6 present these statistics for all POS tags, comparing human and system annotations. In particular, it can be seen how the scores for NOUN were obtained: the precision of 2/3 represents the fraction of correct system predictions, while the recall of 2/4 represents the fraction of human-annotated NOUN tags that were correctly found by the system.⁴

In this thesis, we will also consider a variant of the F_1 score known as the Best F_1 score (BF_1). This measure is useful when comparing continuous system predictions to categorical judgments in a dataset. It is obtained by calculating the F_1 score for the top k entries classified as positive (i.e. the k highest-scoring system predictions), for all possible values of k . In other words, for a dataset X

⁴ To handle edge cases, we define 0/0 as 1.0.

POS tag	Total human	Total system	True Positives	Precision (TP/system)	Recall (TP/human)	F ₁
ADJ	1	1	1	1/1 = 1.00	1/1 = 1.00	1.00
ADP	4	4	4	4/4 = 1.00	4/4 = 1.00	1.00
ADV	0	1	0	0/1 = 0.00	0/0 = 1.00	0.00
DET	2	2	2	2/2 = 1.00	2/2 = 1.00	1.00
NOUN	4	3	2	2/3 = 0.67	2/4 = 0.50	0.57
PROPN	4	3	3	3/3 = 1.00	3/4 = 0.75	0.86
VERB	1	2	1	1/2 = 0.50	1/1 = 1.00	0.67

Table 2.6.: Binary evaluation measures for POS-tag prediction.

ordered by predicted scores:

$$BF_1(X) = \max_k F_1(X_{1...k}).$$

This measure will be used in Chapter 5 when comparing the continuous predictions from our system with the categorical judgments from the *Farahmand* dataset.

2.3. Multiword expressions

Sentences in human languages are more than an unordered collection of words. In order to convey a specific meaning, the words in a sentence must be structured, grouped so as to create phrases, which themselves can be recursively grouped until the whole sentence has been internally connected. The semantics of the whole sentence can then be derived from the semantics of its individual components and from the way in which they relate to each other.

In every known human language, there is a class of expressions that does not necessarily behave in this compositional manner, known as multiword expressions (MWEs). Examples of MWEs would be the English verb–particle construction *give up*, the French NC *carte bleue* (‘bank card’, lit. *blue card*), and the Portuguese idiom *bater as botas* (‘kick the bucket, die’, lit. *hit the boots*). The meaning of an MWE is not always formed by the application of regular rules from the grammar. Rather, each MWE constitutes a semantic unit that spans over multiple lexemes (Sag, Baldwin, Bond, et al. 2002), and which often needs to be analyzed as an indivisible entity. MWEs may present lexical, morphological, syntactic, semantic, pragmatic and statistical idiosyncrasies (Baldwin and Kim 2010), as exemplified below:

- Lexical idiosyncrasy: MWEs may contain words that do not otherwise exist in the language, and thus cannot appear by themselves⁵. Examples include

⁵ Often called *cranberry words*.

EN *of yore*; PT *no entanto* ‘however’ (lit. *in-the entanto*); and FR *au fur et à mesure* ‘in keeping with’ (lit. *to-the fur and to-the measure*).

- Morphological idiosyncrasy: MWEs may contain words that do not respect the normal rules of inflection, as in EN *spill the beans*, where the word *bean* must always be pluralized; PT *azul-marinho* ‘navy blue’ (lit. *marine blue*), where the *marine* adjective does not inflect; and FR *grand-mère* ‘grand-mother’ (lit. *big mother*), where the adjective *big* does not inflect.
- Syntactic idiosyncrasy: MWEs may not conform to the regular syntactic rules of the language. Examples include EN *by and large*, which behaves as an adverb; PT *faz de conta* ‘make-believe’ (lit. *makes of account*), which behaves as a noun; and FR *bon marché* ‘cheap’ (lit. *good market*), which behaves as an adjective. In the case of *by and large*, the MWE comprises a sequence of elements (preposition+conjunction+adjective) that would otherwise not even be allowed by the grammar of the English language.
- Semantic idiosyncrasy: the meaning of the whole expression may not come from the combination of the meaning of its parts. This can be seen in the idiomatic meaning of an MWE such as EN *to kick the bucket*, with the equivalent PT *bater as botas* (lit. *to hit the boots*) and FR *casser sa pipe* (lit. *to break one’s pipe*).
- Pragmatic idiosyncrasy: MWEs may only occur in a particular extra-linguistic context, as in the case of EN *all aboard*, PT *bom dia* ‘good morning’ (lit. *good day*), FR *au revoir* ‘goodbye’ (lit. *to see-again*).
- Statistical idiosyncrasy: the expression may have been conventionalized in a specific form, even though a substitution by a synonym could happen in principle. This is the case of EN *many thanks* (compare with **a lot of thanks*), PT *café preto* ‘black coffee’, compare with **café negro*), FR *noir et blanc* ‘black and white’, compare with **blanc et noir*).

Note that the definition of MWEs in the literature may exclude purely statistically idiosyncratic expressions (also known as collocations). In this thesis, an MWE is understood as a more general term that encompasses all types of idiosyncratic units that cross word boundaries (Sag, Baldwin, Bond, et al. 2002; Baldwin and Kim 2010), as the main interest of our research is the variation of these expressions in a continuum of idiomaticity.

The idiosyncratic behavior of MWEs might lead one to think that such expressions constitute rare exceptions in the language. However, estimations on the number of MWEs in a given speaker’s lexicon may reach at least the same order of magnitude as the number of single words (Jackendoff 1997). In the context of specialized domains, this number is expected to be even higher, as they naturally favor the apparition of multiword technical terms (Sag, Baldwin, Bond, et al. 2002).

The abundance of MWEs, coupled with the fact that their occurrences are not obvious and their meaning is not predictable, contributes to making this an essential topic of research for NLP (Savary, Sailer, Parmentier, et al. 2015). A correct analysis of MWEs is essential to any computer application that has the goal of somehow interpreting written texts, such as machine translation and text simplification.

2.3.1. Nominal compounds

A category of MWE that is of particular interest in this thesis is the *nominal compound*. A nominal compound (NC) is defined as a syntactically well-formed and conventionalized noun phrase containing two or more content words, the head of which is a noun.⁶ Such compounds can express different levels of semantic idiosyncrasy: their interpretation may come directly from the meaning of its components (e.g. *climate change*), or be highly idiomatic (e.g. *cloud nine*), with partially idiomatic cases in-between (e.g. *spelling bee*, *middle school*) (Nakov 2013).

The syntactic realization of NCs varies across languages. In English, they are often expressed as a sequence of nouns, usually $N_1 N_2$ (with the head noun N_2 modified by N_1). This is the most frequent annotated POS-tag pattern in the MWE-annotated English corpus DiMSUM (Schneider, Hovy, Johannsen, et al. 2016). In French and Portuguese, NCs often assume the form of $ADJ N$ or $N ADJ$, where ADJ is an adjective that modifies the head noun N . Examples of these constructions include the French $ADJ N$ compound *petite annonce* ('classified ad', lit. *small announcement*) and the Portuguese $N ADJ$ compound *buraco negro* ('black hole', lit. *hole black*). Additionally, NCs in the three languages may also include prepositions that provide a hint on the role of the modifier noun with respect to the head (e.g. the $N PREP N$ compound *rule of thumb*). Most prepositions are highly polysemous, and it is not clear how they should be represented in the context of distributional semantic models. In the remainder of this thesis, we will focus on 2-word NCs that follow the form $N_1 N_2$ (in English), $N ADJ$ (in Portuguese and French) and $ADJ N$ (in the three languages).

2.3.2. Type discovery

The goal of MWE discovery is to automatically find new MWEs in corpora, collecting these MWEs in a lexicon for future use (Constant, Eryiğit, Monti, et al. 2017). The earliest works toward MWE discovery and lexicon building involved attempts at extracting collocations. These are expressions that present some level of statistical idiosyncrasy, regardless of other levels of idiosyncrasy. By tak-

⁶ The terms *noun compound* and *compound noun* are usually reserved to noun-noun compounds. These are typical of Germanic languages, but not as common in Romance languages.

ing into account some statistical measures, pairs of collocated words could be extracted from corpora with high accuracy (Smadja 1993).

Subsequent works have applied linguistic knowledge in order to target specific categories of MWE, such as noun compounds (e.g. *milk tooth*), verb–particle constructions (e.g. *look up*) and light-verb constructions (e.g. *take a shower*) (Justeson and Katz 1995; Frantzi, Ananiadou, and Mima 2000; Stevenson, Fazly, and North 2004; Evert and Krenn 2005). More recent works also focus on more general type-independent approaches to MWE discovery (Seretan 2011; Agrawal, Aggarwal, et al. 2013; Tsvetkov and Wintner 2014).

MWE discovery techniques usually rely on word frequency and association measures, which are inexpensive language-independent methods of detecting the conventionalization of MWEs. The construction of such MWE lists can be greatly simplified by using an MWE extractor, such as the *mwetoolkit* (Ramisch 2015). This toolkit includes multiple tools for MWE discovery and manipulation, including an extraction algorithm that builds upon the notion of regular-expression patterns to match token properties. For example, given a noun compound pattern such as (Noun Noun⁺) and a POS-tagged corpus, the extraction yields all occurrences of at least two subsequent nouns in the text. The *mwetoolkit* can then be used to calculate association measures and to filter out MWE candidates.

A major downside of such commonly used extraction techniques is that they do not take semantic information into account when filtering MWE candidates. One of the contributions of this thesis is an extension of the *mwetoolkit* that calculates whether an extracted expression can be treated as a combination of its parts or whether it should be treated as a standalone semantic unit (see Chapter 4 for more details).

2.3.3. Token identification

While MWE type discovery focuses on building lexicons of new MWE types, token identification has the goal of automatically annotating the tokens that correspond to an MWE occurrence in running text (Constant, Eryiğit, Monti, et al. 2017). The accurate identification of MWE tokens is a fundamental task in the pipeline of many NLP applications. For example, MT systems need to know when a group of words must be translated as a unit, and parsers need to recognize the cases where a seemingly unrelated group of words should be joined as a single lexeme or constituent.

Identification of MWE tokens in corpora usually requires an MWE lexicon as input, and can be seen as a tagging process, akin to POS tagging. The goal is to look for occurrences of MWEs in a corpus and to output an augmented version of the corpus that explicitly indicates where each expression occurs. This indication can range from simply joining the MWE components as a single word (using a special “MWE separator” character) to more complex metadata representations,

such as indicating each MWE by the index of its component tokens (Schneider, Hovy, Johannsen, et al. 2016; Savary, Ramisch, Cordeiro, Candito, et al. 2017).

MWE identification tools such as jMWE (Kulkarni and Finlayson 2011) are often used to annotate sentences based on preexisting lexicons. Finite-state transducers can also be used to take into account the internal morphology of component words and perform efficient tokenization based on MWE dictionaries (Savary 2009). The problem of MWE identification has also been modeled using supervised machine learning, where the data is encoded in a begin-inside-outside scheme, from which one can learn sequence taggers such as CRFs (Constant and Sigogne 2011; Schneider, Danchik, Dyer, et al. 2014; Scholivet, Ramisch, and Cordeiro 2017).

These solutions have some shortcomings. One of the problems is that MWEs do not always appear contiguously in the text. For example, they may contain internally inserted modifiers or arguments, as in the expressions *to give [something] up* and *to take [a very long] shower*. In such cases, contiguous identifiers will fail to detect these MWEs. One way of dealing with this problem would be the use of parsing-based approaches (Constant and Nivre 2016), but these require the existence of annotated treebanks for training, which are not available for most languages.

Another shortcoming of using separate tools for type discovery and token identification is that one misses the opportunity of sharing information. This has negative results both in terms of CPU time and in the inability to guarantee that all MWE candidates extracted by one tool have been projected back onto the source corpus by the other tool. One of the contributions of this thesis is an extension of the mwetoolkit that identifies MWE occurrences in text. The implementation addresses the latter by allowing the integration of type and token identification in the same pipeline, and the former by enabling non-contiguous matches (e.g. *eat [food] up*, as in *eat that wonderful chocolate cake up*) and optional and variable elements in MWEs (e.g. *throw [person] to the lions/wolves*). See Section 6.1.1 for more details.

MWE token identification is also closely related to the problem of disambiguation in lexical semantics (Schneider, Hovy, Johannsen, et al. 2016). For example, a machine translation system would need to decide whether an expression such as *a piece of cake* should be interpreted as a single unit (e.g. *the test was a piece of cake*) or as a composition of its parts (e.g. *he just ate a piece of cake*), so that it may translate it into an equivalent meaning in the target language. Semantically-aware MWE token identification is a current topic of research (Constant, Eryigit, Monti, et al. 2017).

In this thesis, we will focus on MWE type-level compositionality prediction. The work on token-level identification is seen as a means to an end: collecting candidate MWE tokens so as to be able to calculate type-level semantics.

2.4. Word semantics

In the past years, NLP work on lexical semantics has been shifting from a focus on symbolic representations (often in the form of handcrafted resources) to more numerical techniques that can automatically extract word representations from large bodies of text. We present below an overview of both approaches.

2.4.1. Symbolic representations

One of the earliest formal representation of semantics was the one promoted by Montague during the 1960s and 1970s. This approach relied on the use of formal logic to treat natural language semantics in the same rigorous way as Chomskian grammar would deal with syntax. The resulting Montague Grammar is still today an active area of research in the field of Linguistics (Partee 2014). Words are defined as sets of elements: the noun *house* stands for the set of everything that could be considered a house, while the adjective *yellow* would be seen as the set of all things yellow. A yellow house would then be anything in this intersection (Baroni and Zamparelli 2010). This representation is somewhat limited, as can be seen in the expressions *red watermelon* (inside) versus *green watermelon* (outside).

The meaning of individual words may also be approximated through a set of semantic labels. This is the approach used in SemCor, one of the earliest sense-annotated corpus for the English language. SemCor was constructed by tagging the content words in the Brown corpus based on a set of semantic *sense* labels⁷ (Landes, Leacock, and Tengi 1998). More recently, the STREUSLE corpus of web reviews has been annotated in terms of on noun, verb and preposition *supersenses*. Supersense tags (such as *person*, *location*, and *event*) are a more coarse-grained way of representing the semantics that allows some level of comparability between different words (Ciaramita and Johnson 2003; Schneider, Hwang, Srikumar, et al. 2016).

Another way of representing semantics is through a graph, where each node represents a different meaning⁸, and edges represent some kind of relation between those meanings. For example, a graph of hypernymy relations (i.e. word generalizations) could contain nodes such as DOG, MAMMAL and ANIMAL, with edges connecting DOG→MAMMAL as well as MAMMAL→ANIMAL. Lexical databases have been constructed for some languages, with the most prominent example being the hand-crafted WordNet (Fellbaum 1998), which includes relations such as hypernymy, meronymy (i.e. part-whole), and antonymy (i.e. words with opposite polarity) for the English language. Such databases require dedicated human

⁷ The sense labels correspond to WordNet synsets.

⁸ Words that may refer to the same meaning are grouped in what is known as a synset, so e.g. the words *dog* and *hound* could be part of a synset called DOG.

work for each word in the language, requiring periodic updates to keep in line with the lexical evolution of the language over time.

Automatically-created databases are a current topic of research, and may be a more feasible alternative to manual annotation (especially for languages that dispose of fewer financial resources), at the expense of accuracy (Navigli and Ponzetto 2012). Such resources may be created from parallel texts, determining the semantics of each word based on its possible translations (Ganitkevitch, Van Durme, and Callison-Burch 2013). However, this approach is inherently reliant on the alignment of large amounts of parallel data, which is a scarce resource for most languages.

2.4.2. Numerical representations

While symbolic approaches often require an impractical amount of high-quality data for a given language or domain, numerical solutions such as distributional semantic models (DSMs) tend to be more malleable in their requirements. DSMs use context information (such as the co-occurrence matrix, described in Section 2.2) to represent the meaning of lexical units in the form of numerical vectors. The central idea is that the meaning of a word is naturally learned by human beings based on the contexts where it appears — or, as popularized by Firth (1957), “you shall know a word by the company it keeps”. Distributional models can be built from any large-enough corpus, and do not particularly require any level of data preprocessing.

Consider the target–context matrix in Table 2.7. Each row corresponds to a target word (*fish*, *dog*, *cat*) followed by its context vector. Each dimension in the context vector represents a measure of co-occurrence between the target (e.g. *dog*) and a context that was seen close to this target in a corpus (e.g. *swim*). For example, the target *fish* has a co-occurrence of 1 in the dimension labeled *leg* and a co-occurrence of 8 in the dimension labeled *swim*.

vocabulary		leg	swim
w_1	fish	1	8
w_2	dog	7	3
w_3	cat	9	2

Table 2.7.: Matrix with targets (rows) \times contexts (columns)

The essence of DSMs consists in treating each of these rows as a numerical D -dimensional representation of its target word (with $D = 2$ in this toy example). The assumption is that the meaning of a word can be derived from this vector. For example, the word *cat* would be represented as the vector $\mathbf{v}(w_3) = [9, 2]$. Figure 2.5 presents a visual representation of these vectors. The results exemplify something that is often seen in DSMs: because both words *dog* and *cat* have a

numerical representation that is strongly associated with *leg* and weakly associated with *swim*, their vectors are closer to each other, from which we can infer that *dog* is overall more similar to *cat* than it is to *fish*.




Figure 2.5.: DSM vectors of the target words “*fish*”, “*dog*” and “*cat*”, considering the contexts “*leg*” and “*swim*”.

The similarity between the vector representation of two words can be mathematically measured through the use of a function such as the cosine, which can be defined for two n -dimensional vectors as their normalized dot product:

$$\cos(\mathbf{v}, \mathbf{w}) = \frac{\mathbf{v} \cdot \mathbf{w}}{|\mathbf{v}| \cdot |\mathbf{w}|},$$

where $|\mathbf{v}|$ represents the norm of the vector \mathbf{v} . This definition can be further expanded as:

$$\cos(\mathbf{v}, \mathbf{w}) = \frac{\sum_{i=1}^n \mathbf{v}_i \cdot \mathbf{w}_i}{\sqrt{\sum_{i=1}^n \mathbf{v}_i^2} \cdot \sqrt{\sum_{i=1}^n \mathbf{w}_i^2}},$$

The cosine has a value closer to 1 for vectors that are closer to each other and closer to 0 for vectors that are perpendicular in most dimensions.⁹ In the case of DSMs, this perpendicularity indicates that those vectors do not share many distributional features, and might suggest that they are not semantically related. Other measures could be used for measuring word vector similarity, in particular the euclidean distance. One of the advantages of the cosine is the fact that it is invariant with regards to vector length (which is not usually considered meaningful in a DSM representation).

Formally, distributional semantic models attempt to encode the representation of each word in a vocabulary V as a vector of real numbers $\mathbb{R}^{|V|}$. Tradi-

⁹ The cosine can also be negative, e.g. the vectors point to opposite directions in most dimensions.

tionally, this representation is built by weighting the level of co-occurrence of all pairs of words. This weighting can either be done by counting the number of co-occurrences (Baroni, Dinu, and Kruszewski 2014) or by calculating some measure of the association between target and context, such as the PPMI (Levy, Goldberg, and Dagan 2015). Words are usually deemed to co-occur if they appear in the text under a small fixed-size window of words. Alternatively, some works define co-occurrence based on whether the two words share a syntactic dependency relation (Lin 1998; Padó and Lapata 2007). In all cases, the result is a target–contexts mapping $M = V \times \mathbb{R}^{|V|}$ where many of the context weights are zero (as most word pairs in $V \times V$ will almost never co-occur), and it is thus often implemented as a sparse matrix. A threshold on the number of word pair co-occurrences is often applied to discard low weights.

An alternative to working on a sparse representation is the use of *word embeddings*, in which the vectors are transformed so as to have a significantly smaller number of dimensions.¹⁰ Two solutions are commonly employed in the literature: global contexts and dimensionality reduction. In the case of global contexts, only the top k most frequent words are considered as contexts, producing a $|V| \times k$ matrix (Salehi, Cook, and Baldwin 2014; Padró, Idiart, Villavicencio, et al. 2014). A second alternative is the use of dimensionality reduction techniques. Assuming that all vectors represent data points on a space whose mean is $\mu = 0$, a technique known as single value decomposition (SVD) may be used to transform the matrix rows $M_1..M_n$ in such a way that the largest variance now occurs on $M_{i,1}$ (i.e. maximizing the variance $\sigma_i^2(M_{i,1})$), the second largest variance on $M_{i,2}$, and so on. Only the first k components of each vector are then kept; the rest is discarded. The rationale is that higher variances represent actual structure, while lower variances represent noise in the data.

SVD achieves its results through a matrix factorization technique (Shlens 2014). Formally, it decomposes the matrix $M_{m \times n}$ into three other matrices:

$$M = U_{m \times m} \cdot \Sigma_{m \times n} \cdot V_{n \times n},$$

where U and V specify rotations and Σ is a diagonal matrix which specifies a scaling operation. The product $U\Sigma$ has the aforementioned property in which lower indexes correspond to higher variances. It can be truncated into an $m \times k$ matrix for smaller values of k , effectively obtaining a version of M that has a reduced number of dimensions. Similar methods have also been published focusing on factorizing the logarithm of the co-occurrence matrix (Pennington, Socher, and Manning 2014) and factorizing a matrix of PPMI values (Salle, Villavicencio, and Idiart 2016).

Another word embedding technique is the one adopted by word2vec, which trains a neural network in a task that involves predicting target–context relationships. Two approaches have been proposed: training a network to predict a

¹⁰ Each of these smaller vectors is then called a *word embedding*.

target word given a window of surrounding contexts, known as the Continuous Bag-of-Words (CBOW) model; and training a network to predict likely contexts for a given target, known as the skip-gram model (Mikolov, Sutskever, Chen, et al. 2013). Figure 2.6 presents the general architecture of both approaches, for a window of $2 + 2$ words around each target. In both cases, input words are represented as a one-hot vector in $\{0, 1\}^{|V|}$. Each entry in the input vector is connected to a hidden layer of d neurons, which learns to generate an $\mathbb{R}^{|V|}$ output that predicts¹¹ the one-hot vector of w_i (for CBOW) or its contexts (for skipgram). After training, the input weights from the hidden layer are then taken as the set of d -dimensional word embeddings. In both word2vec approaches, the network automatically adapts itself to encode useful semantic information as a side effect of trying to solve its prediction task.

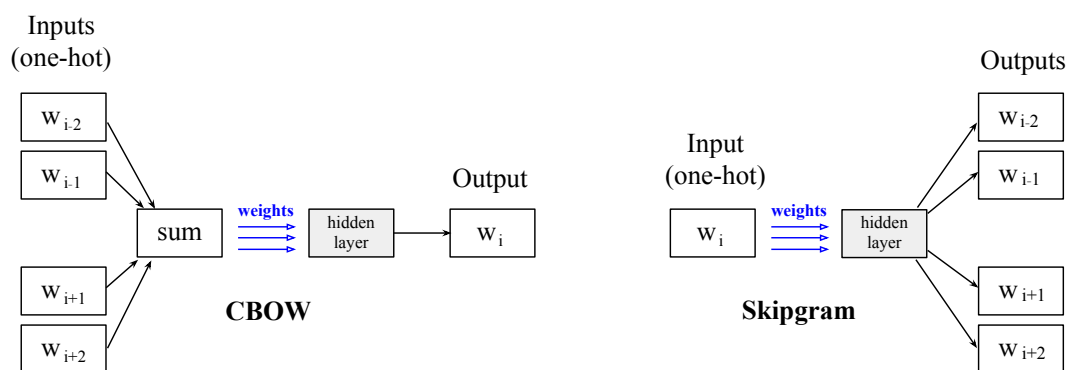


Figure 2.6.: Architecture of word2vec (CBOW and skipgram).

When building a distributional semantic model, there are some common parameters that must be considered, but that are orthogonal to the choice of DSM technique. One of the main considerations is the type of information that will be provided for each word in the corpus: using the word's surface form may generate a sparser model (due to inflections), while using the lemma may merge unrelated occurrences, ignoring relevant morphological distinctions. The use of POS tags may also contribute to disambiguate polysemous words (e.g. compare *the circle* with *we circle*), but it risks the introduction of tagging errors. Another consideration is the removal of stopwords: common function words (such as *the*, *of*, and *for*) do not contribute much to the semantics of the text, and their removal may allow DSM techniques to better capture relevant co-occurrence patterns in the text. In Chapter 5, we present some contributions to the understanding of how these parameters affect different DSMs.

¹¹ During training, the output layer is followed by a softmax layer, which generates the probability for each element in V (and which is compared to the one-hot in the output).

2.5. MWE semantics

Human language is often modeled as though it were compositional, i.e. assuming that the meaning of the whole can be built from the meaning of its parts. However, MWEs may display a wide range of idiomaticity, ranging from compositional cases (e.g. *tennis championship*) to idiomatic non-compositional cases (e.g. *gravy train*). For the latter, the meaning of the expression cannot be understood directly from the meaning of its parts (e.g. a *gravy train* refers to a low-effort lucrative endeavor). Even when there is a level of compositionality in the expression, the contribution of each word may vary considerably, independently from its status as a syntactic head or modifier, as *tears* (head) in *crocodile tears* versus *cash* (modifier) in *cash cow*. In this section, we present the state-of-the-art approaches towards the representation of compositional and idiomatic MWE semantics, using symbolic as well as numerical representations.

2.5.1. Symbolic representation

In the case of compositional MWEs, Lauer (1995) argues that prepositions (such as *from*, *for*, *in*) can be used to classify the role of each component (e.g. *olive oil* is *oil from olives*). These prepositions are explicitly part of some NCs in Romance languages (e.g. FR *huile d'olive* and PT *azeite de oliva*). More generally, Girju, Moldovan, Tatu, et al. (2005) present and compare several inventories of semantic relations between nouns inside NCs, ranging from fine-grained to coarse senses. These relations include syntactic and semantic classes such as *subject*, *instrument* and *location*.

Free paraphrases have also been used to model compositional MWE semantics based on the meaning of the components. Nakov (2008) suggests using unsupervised generation of paraphrases combined with web search engines to classify NCs. This was further extended in SemEval 2013, in a task where free paraphrases were ranked according to their relevance for explicitly describing the underlying semantic relations in the compounds (Hendrickx, Kozareva, Nakov, et al. 2013). For instance, with respect to the expression *flu virus*, the paraphrases at the top of the rank contained verbs such as *cause*, *spread* and *create* (i.e. *virus that causes/spreads/creates the flu*).

Regarding the representation of idiomatic MWEs, lexical resources such as the aforementioned WordNet (Fellbaum 1998) will often include them alongside single words. However, the drawbacks from single-word lexical resources still hold true for MWEs. In particular, for automatically constructed lexicons, the relatively lower frequency and high syntactic variability of MWEs may exacerbate the discrepancy between single-word and MWE coverage even further.

Regarding the representation of idiomatic MWEs, some works extend the *supersense* approach used for single words to identify every MWE as an indivisible unit pertaining to a semantic class. Recent versions of the SemCor corpus (Lan-

des, Leacock, and Teng (1998) annotate MWEs as well as single-word units based on a set of supersenses derived from top-level WordNet hypernym senses. A similar approach is adopted by the STREUSLE corpus, which contains supersense labels for nouns, verbs and prepositions, both when acting as single words and as part of an MWE (Schneider, Onuffer, Kazour, et al. 2014b; Schneider, Hwang, Srikumar, et al. 2016). This latter corpus was also extended in the DiMSUM shared task, with the joint goal of token-based MWE identification and sense disambiguation (Schneider, Hovy, Johannsen, et al. 2016).

2.5.2. Numerical representation

Numerical approaches to MWE semantic representation usually focus on a continuum of compositionality ranging from very compositional expressions to very idiomatic ones. The meaning of an MWE is then expressed through a numerical compositionality score. A low score indicates a completely idiomatic meaning, while a high compositionality score indicates that the meaning of the MWE comes directly from its parts. For example, using a range from 0 to 1, the idiomatic expression *sitting duck* could be associated with the compositionality score 0.2, while the compositional expression *swimming pool* could be assigned a high score such as 0.9.

Separate scores can also be used to represent the literality associated with each individual word (Reddy, McCarthy, and Manandhar 2011). For example, the expression *spelling bee* could be 80% literal with regards to *spelling* and 0% literal with regards to *bee*, while *milk tooth* could be 20% literally related to *milk* and 100% related to *tooth*.

The meaning of an MWE can also be specified in terms of its *conventionalization* (Farahmand, Smith, and Nivre 2015). Just like compositionality scores can distinguish MWEs in a continuum of idiomaticity, conventionalization scores can be used to identify the level of perceived statistical idiosyncrasy in an expression. For example, the expression *tap water* could be considered 20% conventionalized, while *spelling bee* could be judged as 100% conventionalized. Note that conventionalization does not imply idiomaticity (e.g. *time machine* is highly conventionalized while being fairly compositional), and both measures could be used in a single dataset to more precisely specify the semantics of an expression.

Annotating the semantics of MWEs is a considerably hard task, and annotators may disagree on the exact compositionality score. Therefore, scores are often average among multiple annotators. One source of divergence that may be found among annotators is that some datasets do not take polysemy into account, as the authors ask annotators to think about the most common sense of an MWE without providing any context. Some of these datasets address this issue by providing example sentences to attenuate this problem.

Numerical scores can sometimes be considered a more flexible alternative to symbolic representations of semantics. For one, they allow the interpretation of

compositionality in a continuum, which is in line with the perception that the meaning of some MWEs can be more easily guessed than the meaning of the others. Numerical representations can be more readily applied in other numerical contexts, such as DSM representations of semantics.

Numerical representations also allow for fine-grained distinctions that are not possible in a strictly categorical setting. For example, while one may stipulate the categories *idiomatic* and *compositional*, any further attempts at representing partial levels of compositionality would rely on an (implicit) ranking among the categories, tending towards the numerical representations. The possibility of fine-grained distinctions can also be seen as one of the major downsides of numerical scores, as it may introduce uncertainty into the dataset due the subtle differences with which different people see the MWEs. Symbolic representations may also be preferable in the case of highly polysemous MWEs, as the distinction between the multiple senses may not be feasible with bases on commonly studied dimensions such as compositionality or conventionalization.

2.5.3. Compositionality datasets in the literature

We present below a list of relevant datasets representing MWEs alongside human-rated compositionality scores:

- Baldwin and Villavicencio (2002) collected binary type-level judgments for 3 078 English phrasal verbs. Each entry is classified by two experts as either compositional (e.g. *give back*) or idiomatic (e.g. *pull over*), with 14% of the MWEs in the dataset being judged as idiomatic.
- McCarthy, Keller, and Carroll (2003) present a dataset of 116 English verb–particle constructions, annotated with type-level compositionality scores by three native speakers, on a scale ranging from 0 (idiomatic) to 10 (compositional). Five of these were unknown to at least one judge, and were removed from the dataset for their experiments.
- Bannard (2006) collected binary judgments for 160 English verb–particle constructions. In this work, compositionality judgments for each expression were collected from multiple annotators, allowing more fine-grained distinctions in idiomaticity.
- Reddy, McCarthy, and Manandhar (2011) collected judgments for a set of 90 English noun–noun and adjective–noun compounds, in terms of three numerical scores: the compositionality of the compound as a whole and the literal contribution of each of its parts individually, using a scale from 0 to 5. Compounds included in the dataset were selected to balance frequency range and degree of compositionality (low, middle and high). The dataset was built through crowdsourcing, and the final scores are the average of

30 judgments per compound. This dataset will be used in our intrinsic evaluation experiments in Chapter 5, where it will be referred to simply as *Reddy*.

- Gurrutxaga and Alegria (2013) had three experts classify 1200 Basque noun–verb expressions according to one of three possibilities: idiomatic, compositional collocation, or free combination.
- Roller, Schulte im Walde, and Scheible (2013) collected judgments for a set of 244 German noun–noun compounds, each compound with an average of around 30 judgments on a compositionality scale from 1 to 7, obtained through crowdsourcing. The resource was later enriched with feature norms (Roller and Schulte im Walde 2014).
- Farahmand, Smith, and Nivre (2015) had 4 experts annotate 1042 English noun–noun compounds. Each annotator provided binary judgments for every MWE regarding idiomaticity (non-compositionality) and conventionalization (when a particular choice of words has been crystallized as part of the language, even if synonyms would have been understandable). A hard threshold can be applied so that compounds are considered as non-compositional if at least two annotators say so (Yazdani, Farahmand, and Henderson 2015), and the total compositionality score is given by the sum of the 4 binary judgments. This dataset will be used in our intrinsic evaluation experiments in Chapter 5, where it will be referred to as .
- Schulte im Walde, Hättö, Bott, et al. (2016) collected judgments for a set of 868 German noun–noun compounds, with human judgments of compositionality ranging on a scale of 1 to 7. The dataset is also annotated for in-corpus frequency, productivity and ambiguity, and a subset of 180 compounds has been selected so as to be balanced with respect to these variables. The different annotations were performed by the paper authors, linguists, and through crowdsourcing. A similar dataset has been collected for verb–particle constructions (Bott, Khvtisavrishvili, Kisselew, et al. 2016).

Some of these datasets were constructed with binary judgments, while others were constructed with a more malleable representation of idiomaticity, requesting human raters to specify their judgment over a range of possible values. Note, however, that even binary judgments could constitute a numerical dataset. As long as there are enough annotators, the average of the judgments can be taken as a numerical estimate of its perceived idiomaticity.

While the compositionality judgments from the datasets above could be used by themselves as features to semantic tasks such as MWE token identification (described in Section 2.3.3), the size of these datasets may be a limiting factor in the results obtained. On the other hand, these datasets are particularly useful as a way of evaluating the quality of automatic models of compositionality

prediction (which themselves may then be used to predict the compositionality of a much larger set of MWEs). In Chapter 3, we present three new datasets with human annotation of compositionality scores. Chapter 5 then evaluates a framework of compositionality prediction on these new datasets, alongside with the *Reddy* and *Farahmand* presented above.

3. Compositionality datasets

As we have seen in Chapter 2, MWE compositionality can be modeled in terms of the contribution of meaning of each element toward the meaning of the whole. Some numerical datasets have been proposed in the literature, but they are restricted to English and German MWEs. Moreover, in the former language, only one relatively small dataset contains non-categorical compositionality scores.

In this chapter, we describe the construction of three new datasets of human-annotated compositionality scores for nominal compounds (NCs). These datasets are necessary for our evaluations of compositionality prediction models (reported in Chapter 5). The resources encompass: 180 French nominal compounds (*FR-comp*); 180 Brazilian Portuguese nominal compounds (*PT-comp*); an extension of the English-language *Reddy* dataset with 90 additional compounds (*EN-comp*₉₀), for a total of 180 English compounds (*Reddy*⁺⁺).

The work presented in this chapter has also been described in three published papers (Ramisch, Cordeiro, Zilio, et al. 2016; Cordeiro, Ramisch, and Villavicencio 2016a; Wilkens, Zilio, Cordeiro, et al. 2017). For French and Portuguese, this is the first human-rated dataset of nominal compound compositionality.

3.1. Data collection

For each of the 3 target languages (English, French and Portuguese), quantitative measures for the level of compositionality of the nominal compounds in the dataset were collected through crowdsourcing. Non-expert participants were asked to judge each compound in the context of three sentences where the compound displayed the same sense, followed by an evaluation of the degree to which the meaning of the compound is related to the meaning of its individual parts. This follows from the assumption that a fully compositional expression will have an interpretation whose meaning comes from both words (e.g. *lime tree*, which is effectively a *tree of limes*), while a fully idiomatic compound will have a meaning that is unrelated to its components (e.g. *nut case*, which refers to an eccentric person and is not related to *nuts* or *cases*). This work follows the protocol from Reddy, McCarthy, and Manandhar (2011), where the compositionality is explained in terms of the literality of the individual parts. This type of indirect annotation of compositionality is less specialized, and does not require expert linguistic knowledge, while still providing reliable data, as will be shown later.

For each language, data collection involved the following four steps: compound selection; sentence selection; questionnaire design; and questionnaire application.

3.1.1. Compound selection

The initial set of idiomatic and partially compositional candidates was constructed by introspection, independently for each language. This list of compounds was complemented by selecting entries from lists of frequent *adjective+noun* and *noun+noun* pairs. These were automatically extracted through POS-sequence queries using the mwetoolkit (Ramisch 2015). The source corpora were ukWaC (Baroni, Bernardini, Ferraresi, et al. 2009), frWaC and brWaC (Boos, Prestes, and Villavicencio 2014), each containing between 1.5 and 2.5 billion tokens.

We avoided selecting compounds in which the head was not necessarily a noun (e.g. FR *aller simple* ‘one-way ticket’ (lit. *going simple*), as *aller* doubles as the noun *going* and the infinitive of the verb *to go*). We also avoided selecting compounds whose literal sense was very common in the corpus (e.g. EN *low blow*). For PT and FR, we additionally discarded the compounds in which the complement was not an adjective (e.g. PT noun–noun *abelha-rainha* ‘queen bee’ (lit. *bee-queen*)), as these constructions are often seen as exocentric (no head/modifier distinction can be made between the compound elements).

For each language, a balanced set of 60 idiomatic, 60 partially compositional and 60 fully compositional compounds was selected by means of a coarse-grained manual pre-annotation. We eschewed any attempts at selecting equivalent compounds for all three languages. A compound in a given language may correspond to a single word in the other languages, and even when it does translate as another compound, its pattern of POS-tags and its level of compositionality may be widely different.

3.1.2. Sentence selection

For each compound, we selected 3 sentences from the WaC corpus where the compound is used with the same meaning. We sorted them by sentence length, in order to favor shorter sentences, and manually selected 3 examples that satisfy these criteria:

- The occurrence of the compound must have the same meaning in all three sentences.
- Each sentence must contain enough context to enable a clear disambiguation of the compound.
- There must be enough inter-sentence variability, so as to provide a higher amount of disambiguating contexts.

The goal of these sentences was to be used as disambiguating context for the annotators. For example, for the compound *benign tumour*, we present the following disambiguating sentences: (1) “Prince came onboard to have a large

benign tumor removed from his head”; (2) “We were told at that time it was a slow growing *benign tumor* and to watch and wait”; (3) “Completely *benign tumor* is oncocytoma (it represents about 5 % of all kidney tumors)”.

3.1.3. Questionnaire design

The questionnaires were presented as online webpages, and followed the same structure for each compound. The questionnaire starts with a set of instructions that briefly describe the task and direct participants to fill an external identification form. This form collects demographics about the annotators, and ensures that they are native speakers of the target language, following Reddy, McCarthy, and Manandhar (2011). This form also presents some example questions with annotated answers for training. After filling in the identification form, users could start working on the task itself. The questionnaire was structured in 5 subtasks, presented to the annotators through these instructions:

1. Read the compound itself.
2. Read 3 sentences containing the compound.
3. Provide 2 to 3 synonym expressions for the target compound seen in the sentences, preferably involving one of the words in the compound. We ask annotators to prioritize short expressions, with 1 to 3 words each, and to try to include the words from the nominal compound in their reply (eliciting a paraphrase).
4. Using a Likert scale from 0 (completely disagree) to 5 (completely agree), judge how much of the meaning of the compound comes from modifier and head separately. Figure 3.1 shows an example for the judgment of the literality of the head (*benign*) in the compound *benign tumor*.
5. Using a Likert scale from 0 (completely disagree) to 5 (completely agree), judge how much of the meaning of the compound comes from both of its components (head and modifier). This judgment is requested through a question that paraphrases the compound: “would you say that a *benign tumor* is always literally a *tumor* that is *benign*?”.

We have been consciously careful about requiring answers in an even-numbered scale (0–5 makes for 6 reply categories), as otherwise, undecided annotators could be biased towards the middle score. As an additional help for the annotators, when the mouse hovers over a reply to a multiple-choice question, we present a guiding tooltip, as in Figure 3.1. We avoid incomplete replies by making Subtasks 3–5 mandatory.

The order of subtasks has also been taken into account. During a pilot test, we found that presenting the multiple-choice questions (Subtasks 4–5) before

4. In your opinion, is a *benign tumor* always literally a *tumor*?

NO 0 1 2 3 4 5 YES

No — I see only a vague relation between a *benign tumor* and a *tumor*

Figure 3.1.: Excerpt from the questionnaire of the compound *benign tumor*, evaluating compositionality regarding the head of the compound.

asking for synonyms (Subtask 3) yielded lower agreement, as users were often less self-consistent in the multiple-choice questions (e.g. replying that “*benign tumor* is not a *tumor*” in Subtask 4 while replying that “*benign tumor* is a *tumor* that is *benign*” in Subtask 5). This behavior was observed even when they later carefully selected their synonyms. Asking for synonyms in Subtask 3, prior to the multiple-choice questions, prompts the user focus on the target meaning for the compound and also have more examples (the synonyms) when considering the semantic contribution of each element of the compound. In this work, the synonyms were only used to motivate annotators to think about the meaning of the compound. In the future, this information could be exploited for compositionality prediction, but also for lexical substitution tasks (Wilkins, Zilio, Cordeiro, et al. 2017).

3.1.4. Judgment collection

Annotators participated via online questionnaires, with one webpage per compound. For EN and FR, annotators were recruited and paid through Amazon Mechanical Turk (AMT). For PT, we developed a standalone web interface that simulates AMT, as Portuguese speakers were rare in that platform. Annotators for PT were undergraduate and graduate students of Computer Science, Linguistics and Psychology. For each compound, we have collected judgments from around 15 annotators.¹

For each compound, the response from all annotators were gathered up into an average compound score. We obtained the following variables:

- c_H : The contribution of the *head* to the meaning of the compound (e.g. is a *busy bee* literally a *bee*?), with standard deviation σ_H .
- c_M : The contribution of the *modifier* to the meaning of the compound (e.g. is a *busy bee* literally *busy*?), with standard deviation σ_M .
- c_{WC} : The degree to which the *whole compound* can be interpreted as a combination of its parts (e.g. is a *busy bee* a *bee* that is *busy*?), with standard

¹ EN includes the 90 compounds from Reddy, McCarthy, and Manandhar (2011), which are compatible with the other 90 compounds collected for the dataset.

deviation σ_{wc} .

The average c scores provide absolute judgments on the compositionality of a compound, ranging from 0 (non-literal or idiomatic) to 5 (literal or compositional). All datasets are freely available online.² For a complete list of all compounds, along with their translation, glosses and collected compositionality scores, we refer to Appendix A (*EN-comp₉₀*), Appendix B (*FR-comp*), and Appendix C (*PT-comp*).

3.2. Dataset analysis

In this section, we analyze some properties of the datasets. These are performed on the 180 compounds of each language. We also present some graphics focusing on the 90 compounds collected in *EN-comp₉₀*, for comparison against the whole dataset of 180 compounds in *Reddy⁺⁺*. In some cases, in order to perform a cross-language analysis of the data, we group the compositionality scores of the three datasets into a single dataset *ALL-comp* with all 3×180 compounds.

3.2.1. Score distribution

Figure 3.2 presents the average scores for the 180 compounds for each language. Each of the 4 graphs is ordered in the x-axis based on the whole-compound compositionality scores (rank-based c_{wc}). Values in the y-axis then present the average score of each compound (value-based c_H , c_M and c_{wc}). The average human judgments confirm that the three datasets are balanced in terms of compound idiomaticity, with the whole-compound scores growing at a mostly linear rate (the correlation between c_{wc} and the list of numbers 1..180 is statistically significant, with Pearson $r > 0.99$ for all four graphs). Moreover, there seems to be a greater agreement between the score for the compound and that of its head/modifier for the two extremes (totally idiomatic and fully compositional), with a greater dispersion of head/modifier scores for partially idiomatic compounds.

3.2.2. Difficulty of annotation

For each compound, the difficulty of annotation can be estimated as the standard deviation (σ) among the compositionality scores provided by multiple human raters. Ideally, if all annotators agreed on a compositionality score, σ should be low. Following Reddy, McCarthy, and Manandhar (2011), we calculated for each language the number of compounds that had standard deviation greater than 1.5. The results are shown in Table 3.1. The largest deviations happened for modifiers, which suggests that adjectives may be harder for humans to judge

² <http://pageperso.lif.univ-mrs.fr/~carlos.ramisch/?page=downloads/compounds>

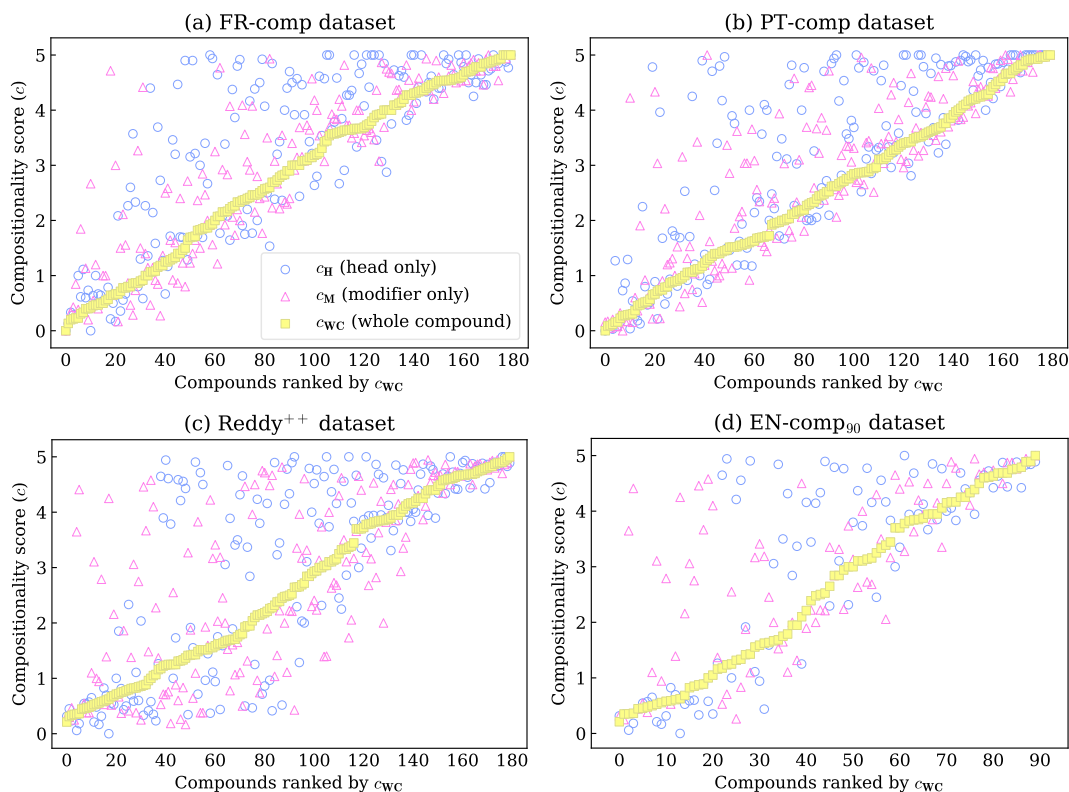


Figure 3.2.: Average compositionality (c_H , c_M and c_{WC}) per compound.

	EN	FR	PT
Compounds with $\sigma_{WC} > 1.5$	22	41	30
Compounds with $\sigma_H > 1.5$	23	44	33
Compounds with $\sigma_M > 1.5$	35	55	34

Table 3.1.: Number of cases of high standard-deviation σ .

than nouns. Indeed, if we consider the average of all standard deviations in *Reddy++*, we obtain $\overline{\sigma_H} = 0.97$ and noun-based $\overline{\sigma_{M=\text{noun}}} = 0.97$ (with 132 cases), but adjective-based $\overline{\sigma_{M=\text{adj}}} = 1.30$ (with 48 cases). This is in line with the average standard deviation found for the other two languages, where every modifier in the dataset is an adjective. For *FR-comp*, $\overline{\sigma_H} = 1.01$ and $\overline{\sigma_M} = 1.18$; while for *PT-comp*, $\overline{\sigma_H} = 0.84$ and $\overline{\sigma_M} = 0.98$.

Figure 3.3 presents the standard deviation scores of every compound as a function of its average compositionality score. Just as the head/modifier-only scores were closer in value to the whole-compound score in the extremities (e.g. highly idiomatic cases with compositionality $c_{WC} < 20$ and highly compositional cases with $c_{WC} > 160$), so are all standard deviations lower in these extremes. This phenomenon may be related to purely statistical effects of extremity values, or may indicate that more extreme judgments are easier for humans to produce. One

consistent property seen among all datasets is that the peak of standard deviation occurs in the left side of the graphs (in particular for σ_H and σ_M), suggesting that idiomatic compounds are slightly harder for humans to judge consistently.

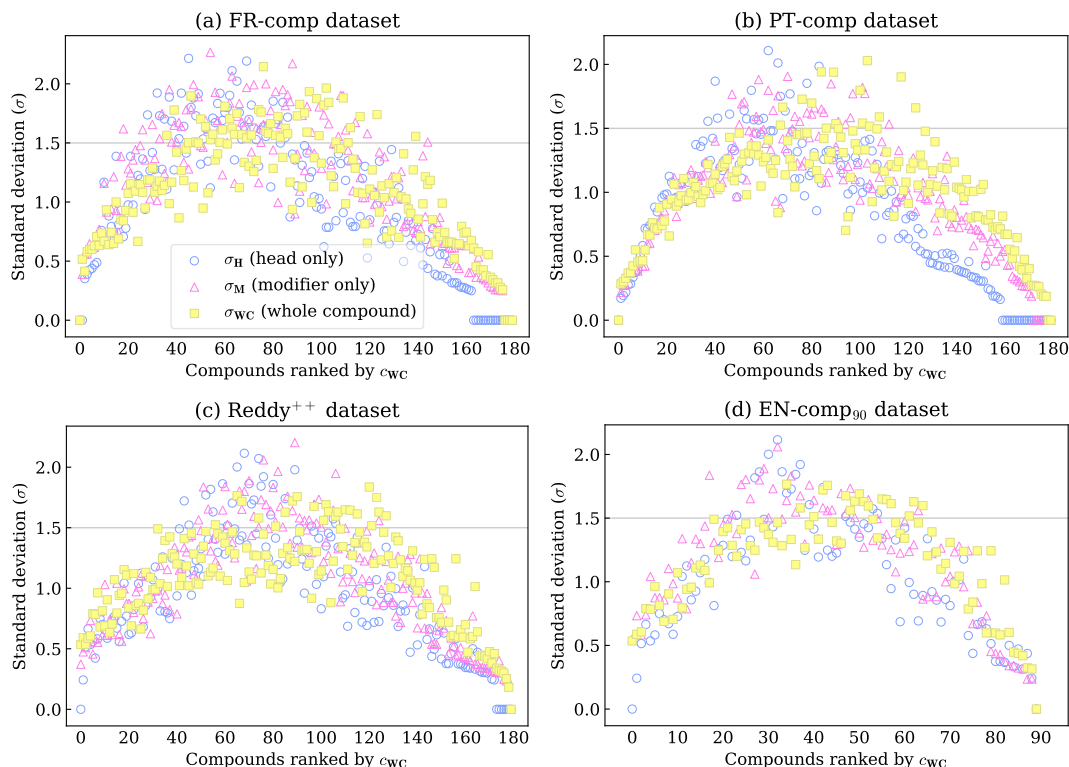


Figure 3.3.: Standard deviation (σ_H , σ_M and σ_{WC}) per compound.

The difficulty of annotation can also be measured through inter-rater agreement measures (described in Section 2.2.4). For the English and French datasets, most participants only provided a small amount of annotations, making these measures unfeasible. For the Portuguese dataset, 3 of the participants annotated a large subset of 119 compounds. For this subset, the pairwise kappa values range from $\kappa = .28$ to $\kappa = .58$ depending on the question (head-only, modifier-only or whole-compound) and on the annotator pair. In the case of α , there was an agreement of $\alpha = .52$ for head-only, $\alpha = .36$ for modifier-only and $\alpha = .42$ for whole-compound compositionality scores. We also calculated the α between an expert annotator and himself some weeks later. The agreement rate ranges from $\alpha = .59$ for whole-compound and modifier-only, to $\alpha = .69$ for head-only compositionality scores.

3.2.3. Estimating whole-compound from head/modifier

A careful analysis of the plot presented in Figure 3.2 suggests that the whole-compound score is lower-bounded by the head-only and modifier-only scores, i.e.

$c_{WC} \approx \min(c_H, c_M)$. This would mean that the whole-compound compositionality score is estimated by human raters based on the literality of its elements. We thus evaluated if it was possible to predict the compositionality score of the whole compound from the scores of its parts. To quantify this relation, we used two models: the arithmetic and geometric mean of the head-only and modifier-only scores for that compound.

Figure 3.4 shows the linear regression for both measures in the *Reddy*⁺⁺ and *FR-comp* datasets. The goodness of fit results for *Reddy*⁺⁺ were $r_{arith}^2 = .90$ for the arithmetic mean, and $r_{geom}^2 = .96$ for the geometric mean, with the latter being a better predictor of the whole-compound compositionality. Similar results were achieved for *FR-comp* ($r_{arith}^2 = .93$, $r_{geom}^2 = .96$), for *PT-comp* ($r_{arith}^2 = .91$, $r_{geom}^2 = .96$) and for *EN-comp*₉₀ ($r_{arith}^2 = .90$, $r_{geom}^2 = .96$). This means that, whenever annotators judged an element of the compound as highly non-literal, they have also rated the whole compound as highly idiomatic. Estimating based on the min operation itself yields similar R^2 values as r_{geom}^2 for *PT-comp* and *FR-comp*. For *Reddy*⁺⁺, $r_{min}^2 = .90$, indicating that the geometric mean is actually a better estimator than the min function itself. The results of this analysis inspired the proposal of the *geom* compositionality prediction strategy, described in Chapter 4.

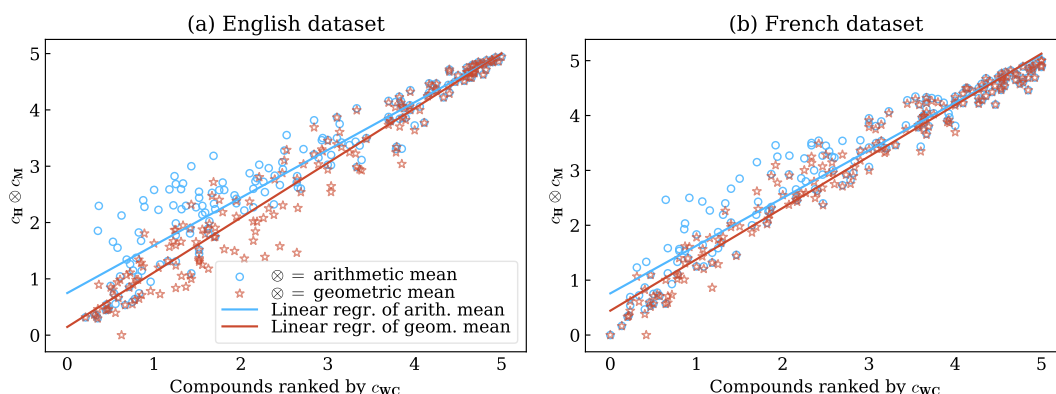


Figure 3.4.: Distribution of $c_H \otimes c_M$ according to c_{WC} of each compound.

3.2.4. Correlation with distributional variables

The hypothesis $h_{idiom} \approx \text{distr.freq}$ suggests that idiomatic MWEs should occur more frequently than compositional ones in general human communication. As a result, there should be a negative correlation between the compositionality score and the frequency of each compound in a sufficiently general corpus. We thus calculate the correlation between the compositionality score c_{WC} of each compound and its frequency in the WaC corpora³. The result is a statistically significant

³ We used the same WaC corpora as in Section 3.1.1.

Spearman correlation of $\rho = .46$ for *Reddy*⁺⁺ and $\rho = .60$ for *FR-comp* (with p-value $p < 10^{-10}$ in both cases). In the case of *PT-comp*, no significant correlation was found (p-value $p > 0.1$). Figure 3.5 presents each dataset where compounds were ordered by frequency and grouped into 18 bins of 10 compounds each. The height of each bar indicates the average of the c_{wc} score assigned by humans to the 10 compounds in the bin. We can see that, in the case of *FR-comp* and *Reddy*⁺⁺, the compounds that are more frequent tend to be assigned higher compositionality scores by humans. These results go against the hypothesis $h_{\text{idiom}} \approx \text{distr.freq}$, which proposed that idiomatic compounds should be overall rather frequent, so as to permit the assimilation of their meaning (Pinker 1995)

Figure 3.5 also presents a graph where the three datasets were combined so as to form a single set of 3×180 compounds. The height of each bin indicates the average score assigned by humans to the 30 compounds therein. As in the case of English and French data above, this combined dataset presents a statistically significant positive correlation of $\rho = .41$ (with $p < 10^{-24}$) between compositionality scores and corpus frequency. This correlation does not mean that all of these idiomatic compounds are rare in an absolute sense, but it does mean that, among the most frequent compounds in the datasets, the majority of entries is more compositional than the average.

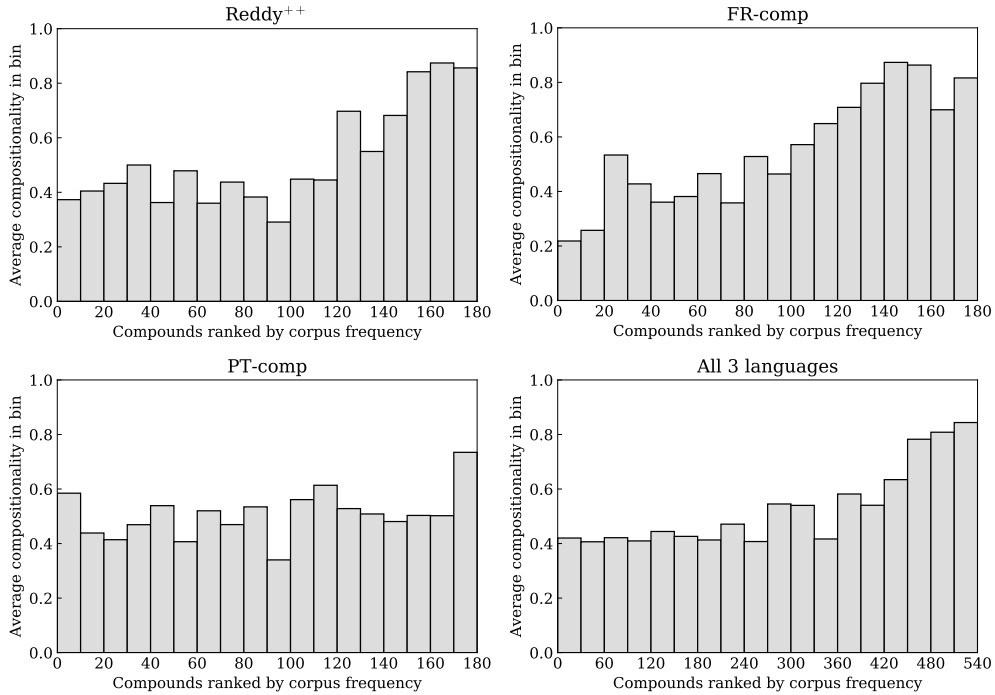


Figure 3.5.: Compositionality for compounds under different frequency bins.

We have similarly analyzed the correlation between the compositionality score of each expression and the level of conventionalization (estimated through the PMI). According to the $h_{\text{idiom}} \approx \text{distr.convent}$ hypothesis, the level of idiomaticity of an

MWE should be positively correlated with the PMI. Figure 3.6 presents all 3×180 compounds ordered by PMI and grouped under 18 bins of 30 compounds each. The height of each bin indicates the average score assigned by humans to the 30 compounds therein. As can be seen, there is no clear pattern of correlation between the two variables. Indeed, differently from the case of the frequency, we found no statistically significant correlation between the compositionality scores and the PMI (this holds true for each dataset by itself, as well as when the 3 datasets are combined). This stands in contrast with the fact that many works in the literature rely on association measures as estimators for compositionality (e.g. using PMI in the discovery of idioms) (Fazly and Stevenson 2006; Bu, Zhu, and Li 2010; Gurrutxaga and Alegria 2013; El Maarouf and Oakes 2015). Given the lack of correlation between these two variables, we do not recommend the use of PMI as an estimator of compositionality.⁴




Figure 3.6.: Compositionality for noun compounds under different PMI bins.

3.3. Summary

In this chapter, we have presented a multilingual group of datasets containing human judgments about the compositionality of nominal compounds. It contains 180 compounds for each of the 3 target languages: English, French and Portuguese. Annotations were collected through crowdsourcing. Since the task was performed by native speakers who may not have a background in linguistics, it needs to be appropriately constrained not to require expert knowledge, and this section has described the methodology that were employed towards that goal.

An analysis of the resulting resource confirmed that the compounds are uniformly distributed across different ranges of compositionality scores. The 3 datasets are comparable regarding their difficulty of annotation, with partially

⁴ We leave the investigation of other association measures for future work.

compositional and adjective-based compounds consistently posing a higher level of difficulty for the human raters. Head-only and modifier-only judgments have also been compared to whole-compound compositionality judgments, with the latter scores behaving as the geometric mean of the former two. Compositionality scores showed no correlation with a measure of conventionalization, differently from what was predicted by the $h_{\text{idiom}} \approx \text{distr.convent}$ hypothesis. Finally, compositionality scores consistently showed a positive correlation with compound frequency in a corpus, refuting the common intuition that the most frequent compounds tend to be idiomatic (hypothesis $h_{\text{idiom}} \approx \text{distr.freq}$).

The datasets presented in this chapter can be used to evaluate applications and tasks requiring some degree of semantic processing, such as lexical substitution and text simplification. For the cases where the numerical judgments alone are not enough for a given task, the datasets also provide sets of paraphrases, which serve as a symbolic counterpart to those scores. These datasets have also been described and analysed in dedicated publications (Ramisch, Cordeiro, Zilio, et al. 2016; Cordeiro, Ramisch, and Villavicencio 2016a; Wilkens, Zilio, Cordeiro, et al. 2017). The complete resource is freely available online.⁵

⁵ <http://pageperso.lif.univ-mrs.fr/~carlos.ramisch/?page=downloads/compounds>

4. Compositionality prediction

Multiword expressions exist in a wide spectrum of idiomaticity, ranging from mere statistically-idiosyncratic compositional combinations of words (such as *beach towel*, which refers to an actual towel), to completely opaque idioms (such as *eager beaver*, which refers to an enthusiastic person). Precision-oriented NLP systems must distinguish between the different levels of compositionality in order to appropriately handle these different kinds of MWEs. The level of MWE compositionality has been measured through a numerical representation in multiple datasets. However, the coverage of these datasets is limited by the availability of human resources.

This chapter focuses on the task of *compositionality prediction*, which consists in automatically identifying the level of compositionality of MWEs without the input of human raters. The core of this thesis consists in the evaluation of a compositionality prediction model under multiple DSMs and with a variety of parameters. Section 4.1 presents the related work on compositionality prediction. Section 4.2 then presents the compositionality prediction model that was proposed and implemented for this thesis. The remainder of this chapter describes the organization of the experiments for the evaluation of this model: corpus pre-processing (Section 4.3), DSMs (Section 4.4), parameters (Section 4.5), and the evaluation setup (Section 4.6). The evaluation of compositionality prediction models can be performed intrinsically or extrinsically:

- Intrinsic evaluation requires the existence of a dataset in which each MWE is associated with a compositionality score (e.g. the datasets presented in Section 2.5, as well as the resources developed in Section 3), serving as a gold standard. The compositionality prediction model is used to predict those scores, which are then directly compared to the gold standard using a correlation measure (such as those described in Section 2.2.5). This is the approach followed in Chapter 5, which analyses the results of compositionality prediction for nominal compounds under thousands of experimental setups.
- In extrinsic evaluation, predicted compositionality scores can be used to decide how an MWE should be treated in NLP systems. For example, in an application such as machine translation, idiomatic MWEs should be identified and translated as an atomic unit. As a consequence, an evaluation of machine translation quality focusing on MWEs would indirectly reflect the ability of the system to predict compositionality (Cap, Nirmal, Weller, et al. 2015; Stymne, Cancedda, and Ahrenberg 2013; Salehi, Mathur, Cook, et al. 2015). A less application-focused alternative would be the evaluation of the usefulness of predicted compositionality scores in a task of identify-

ing idiomatic MWE occurrences in a corpus. The predicted compositionality scores would then be used as a feature in the underlying MWE identification system. The evaluation of this system would then compare the automatically identified MWEs with a gold standard corpus, resulting in an indirect evaluation of the compositionality prediction scores. This is the approach followed in Chapter 6.

4.1. Related work

Compositionality prediction techniques usually involve measuring the extent to which the meaning of an expression is constructed from a combination of the meaning of its parts. One of the most common setups requires three ingredients: (1) vector representations of single word meanings, such as those built using DSMs; (2) a mathematical model of how the compositional meaning of a phrase should be calculated, as a combination of the single-word meaning of its parts; and (3) a measure of similarity, used to compare the compositionally-constructed meaning of a phrase and its own meaning derived from corpora.¹ Figure 4.1 presents this compositionality prediction architecture, along with the three main ingredients. For each MWE (e.g. *flea market*), the DSM vectors of its elements are

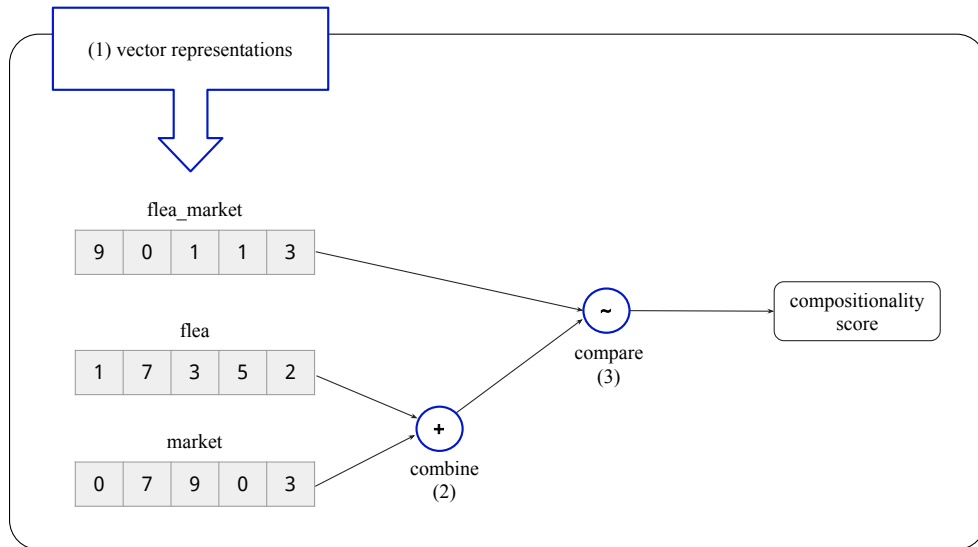


Figure 4.1.: Common compositionality prediction architecture

combined into a single vector, which is then compared against the vector for the MWE built from its occurrences in the corpus. For each of the three ingredients, there are a number of different alternatives that can be seen employed in the

¹ These setups often assume that each word corresponds to a single meaning (i.e. no ambiguity is taken into account).

literature. Throughout this thesis, we will refer to a specific choice of the three ingredients as a *compositionality prediction model*.

The 1st ingredient applies to any kind of numerical representation of word semantics. All of the representations discussed in Section 2.4.2 can be equivalently used in this architecture, and different works in the literature will focus on different representations (e.g. Reddy, McCarthy, and Manandhar (2011) use a co-occurrence matrix with a limited vocabulary V comprising $|V| = 10\,000$ words, while Yazdani, Farahmand, and Henderson (2015) use word embeddings with a reduced number of dimensions). In most works, only a single DSM system is evaluated, under a limited set of parameters. This thesis will consider multiple DSMs under a variety of corpus and DSM parameters.

The 2nd ingredient concerns the mathematical model of meaning combination. One of the most natural choices is the additive model, in which the compositional meaning of an MWE $w_1 w_2 \dots w_N$ is predicted as a linear combination of the word vectors of its components: $\sum_i \beta_i \mathbf{v}(w_i)$, where the β coefficients assign different weights for the representation of each word, and $\mathbf{v}(w_i)$ is a D -dimensional word vector for word w_i , with $D \leq |V|$ (Reddy, McCarthy, and Manandhar 2011; Schulte im Walde, Müller, and Roller 2013; Salehi, Cook, and Baldwin 2015). These different weights can capture asymmetric contributions by each of the components (Bannard, Baldwin, and Lascarides 2003; Reddy, McCarthy, and Manandhar 2011). For example, in the expression *couch potato* (which refers to a stereotypical person who spends a lot of time sitting down and watching television), it is the first word that has a clear contribution to the word meaning, and the highest weight should be in β_1 . In the MWE *flea market*, it is the second word that contributes the most, and the highest weight should thus be in β_2 .

The additive model of composition can be generalized so as to use a matrix of multiplicative coefficients, which can be estimated through linear regression (Guevara 2011). This model can be further modified so as to learn polynomial projections of higher degree, with quadratic projections yielding particularly promising results (Yazdani, Farahmand, and Henderson 2015). These models come with the caveat of being supervised approaches, thus requiring some amount of pre-annotated data in the target language. Due to these requirements, most works focus on unsupervised compositionality prediction methods only, based exclusively on large monolingual unannotated corpora. The latter is also the approach adopted in this thesis.

Alternatives to the linear models include the multiplicative model and its variants (Mitchell and Lapata 2008). However, results suggest that this representation yields inferior results when compared to the predictions obtained through the additive model (Reddy, McCarthy, and Manandhar 2011; Salehi, Cook, and Baldwin 2015). Recent work on predicting intra-MWE semantics also supports the hypothesis that additive models tend to yield better results (Hartung, Kaup-

mann, Jebbara, et al. 2017). This thesis evaluates different variants of compositionality prediction models (see the prediction strategies in Section 4.2).

The 3rd ingredient is a measure of similarity, used to compare the MWE with the sum of its parts. Most works in the literature rely on cosine similarity (Schone and Jurafsky 2001; Mitchell and Lapata 2008; Farahmand, Smith, and Nivre 2015) but it also can be calculated in terms of the overlap between the profiles of word distribution (McCarthy, Keller, and Carroll 2003), assuming that compositional expressions are more similar or share more semantic neighbors with their components than idiomatic ones. In this thesis, the cosine similarity will be used for all evaluations.

Compositionality prediction can also be achieved through an estimation of the likelihood of an MWE (e.g. *red-blood cell*) being replaced by single-word terms in a corpus (e.g. *erythrocyte*). MWEs that can be replaced by many single-word terms are then deemed idiomatic (Riedl and Biemann 2015). An alternative method would be to compare an MWE and its constituents across multiple translations of a text. If the MWE is translated literally, it is predicted as compositional, while non-literal translations are interpreted as an indication of idiomaticity (Salehi, Cook, and Baldwin 2014). The number of possible translations has also been used as an indicator of idiomaticity (Cap 2017).

The level of compositionality of MWEs may also be predicted in the context of MWE-annotated sentences. This is particularly beneficial in the presence of ambiguous MWEs, whose degree of compositionality depends on the context (Sporleder and Li 2009). One such type of ambiguity arises from the possibility of literal and non-literal interpretations for the same lexical unit (e.g. in the expression *spill the beans*). For instance, Köper and Walde (2016) evaluate the impact of different features on the prediction of the literality of German verb-particle constructions. Their features range from the use of a bag-of-words model to distributional statistical scores. The experiments in this thesis focus on non-ambiguous MWEs, so we do not present results for context-dependent compositionality prediction.

4.2. Proposed model

The compositionality principle assumes that the meaning of phrases and sentences can be derived from a combination of the meaning of their components. While this may hold for compositional MWEs, for idiomatic cases we expect the opposite to be true: by combining the semantic representations of the parts of an MWE, we should obtain a representation that is *different* from the representation of the MWE derived directly from corpora. This behavior can be exploited in the construction of the *compositionality prediction model* that was implemented for this thesis and which we evaluate in Chapters 5 and 6.

For each MWE (e.g. *flea market*), let the unitized representation \mathbf{v}_u be the vector representation built for the MWE as a whole, as seen in the corpus²:

$$\mathbf{v}_u(w_1 w_2 \dots w_N) = \mathbf{v}(w_1_w_2_ \dots_ w_N).$$

Define the combined vector \mathbf{v}_β as a function of the individual meaning of the MWE elements (e.g. *flea* and *market*). This vector is calculated through an additive operation of vector composition³:

$$\mathbf{v}_\beta(w_1 w_2 \dots w_N) = \beta_1 \frac{\mathbf{v}(w_1)}{\|\mathbf{v}(w_1)\|} + \beta_2 \frac{\mathbf{v}(w_2)}{\|\mathbf{v}(w_2)\|} + \dots + \beta_N \frac{\mathbf{v}(w_N)}{\|\mathbf{v}(w_N)\|}.$$

where $\beta_i \in [0, 1]$ is a parameter that controls the relative importance of each word for the combined representation, with $\sum_{i=1}^N \beta_i = 1$. The compositionality prediction model may then calculate the compositionality score CS as the cosine similarity between the unitized representation \mathbf{v}_u and the combined representation \mathbf{v}_β :

$$\text{CS}_\beta(w_1 w_2 \dots w_N) = \cos(\mathbf{v}_u(w_1 w_2 \dots w_N), \mathbf{v}_\beta(w_1 w_2 \dots w_N)).$$

The above definition of the compositionality score leaves the precise values of β unspecified. In this thesis, we will consider the following composition strategies:

- Uniform, which defines $\beta_i = \frac{1}{N}$ (e.g. $\beta_1 = \beta_2 = 0.5$ for a 2-word MWE). Equal weights are assigned to every word in the MWE, assuming that they all contribute equally to the meaning of the MWE (as in the MWE *access road*). This is the most commonly used prediction strategy in the literature (Mitchell and Lapata 2010), and most of the results in this thesis will focus on the scores obtained through this weighting strategy.
- Head, for 2-word MWEs, which defines $\beta_{\text{head}} = 1$ and $\beta_{\text{mod}} = 0$, i.e. the modifier is considered to make no contribution to the semantics of the MWE, and the meaning comes from the head alone (as in *crocodile tears*).
- Mod, for 2-word MWEs, which defines $\beta_{\text{head}} = 0$ and $\beta_{\text{mod}} = 1$, i.e. the meaning of the MWE is assumed to come from the modifier, while the head is assumed to make no contribution towards the meaning of the whole (as in the expression *night owl*).

² In the corpus preprocessing stage, the components of the target MWEs are linked by “_” to be treated as a single token; e.g. $w_1_w_2 = \text{flea_market}$ (see Section 4.3).

³ DSM vector length is usually not considered meaningful, and is implicitly normalized during the calculation of the cosine. We do explicitly normalize vector length for the computation of \mathbf{v}_β to properly apply the β weights.

- Maxsim, where the set of weights β_i in the construction of \mathbf{v}_β is defined so as to maximize the value of CS, i.e.:

$$\beta = \operatorname{argmax}_X \text{CS}_{\beta=X}(w_1 w_2 \dots w_N).$$

As a consequence, *maxsim* is capable of expressing the 3 previous prediction strategies as a combination of weights. This model has been developed for the purpose of this thesis. The underlying hypothesis ($h_{\text{strat.maxsim}}$) is that this model is a better predictor of compositionality scores for compositional MWEs, as it constructs a vector with weights that are optimal for a compositional reading.

Note that, in the case of $N = 2$, a closed formula can be derived for the calculation of β values, which avoids an exhaustive search of the parameter space (and is what we used for the experiments in Chapter 5). Let $\beta_2 = 1 - \beta_1$. We want to perform the following maximization:

$$\beta_1 = \operatorname{argmax}_y \cos \left(\mathbf{v}_u(w_1 w_2), y \frac{\mathbf{v}(w_1)}{\|\mathbf{v}(w_1)\|} + (1 - y) \frac{\mathbf{v}(w_2)}{\|\mathbf{v}(w_2)\|} \right).$$

This can be achieved by differentiating the right side of the equation:

$$\frac{d}{d\beta_1} \cos \left(\mathbf{v}_u(w_1 w_2), \beta_1 \frac{\mathbf{v}(w_1)}{\|\mathbf{v}(w_1)\|} + (1 - \beta_1) \frac{\mathbf{v}(w_2)}{\|\mathbf{v}(w_2)\|} \right) = 0.$$

Replacing the cosine by the definition based on the dot product (see page 45) and solving for β_1 , we obtain the closed-form solution:

$$\beta_1 = \frac{\cos(\mathbf{v}(w_1), \mathbf{v}_u(w_1 w_2)) - \cos(\mathbf{v}(w_1), \mathbf{v}(w_2)) \cdot \cos(\mathbf{v}(w_2), \mathbf{v}_u(w_1 w_2))}{(1 - \cos(\mathbf{v}(w_1), \mathbf{v}(w_2))) \cdot (\cos(\mathbf{v}(w_1), \mathbf{v}_u(w_1 w_2)) + \cos(\mathbf{v}(w_2), \mathbf{v}_u(w_1 w_2)))}.$$

We additionally consider two variations on the compositionality score that do not rely on the construction of a combined representation \mathbf{v}_β . They compare \mathbf{v}_u directly to the individual word vectors instead. The two relevant composition strategies are:

- Arith, which calculates the cosine between the MWE and each component individually, and yields the arithmetic mean of the cosines as the compositionality score; i.e.:

$$\text{CS}_A(w_1 w_2 \dots w_N) = \frac{1}{N} \left(\sum_{i=1}^N \cos(\mathbf{v}_u(w_1 w_2 \dots w_N), \mathbf{v}(w_i)) \right).$$

While some works in the literature seem to favor the *uniform* model, some results have been published for *arith* (Reddy, McCarthy, and Manandhar

2011; Schulte im Walde, Müller, and Roller 2013; Salehi, Cook, and Baldwin 2015). To date, no work has compared the behavior of these two models. Given that both strategies represent an additive interpretation of the vectors, our hypothesis ($h_{\text{strat.arith} \approx \text{strat.uniform}}$) is that the highest-ranking configurations of both models should obtain similar scores.

- **Geom**, which calculates the cosine between the MWE and each component individually, and yields the geometric mean of the cosines as the compositionality score; i.e.:

$$\text{CS}_G(w_1 w_2 \dots w_N) = \left(\prod_{i=1}^N \cos(\mathbf{v}_u(w_1 w_2 \dots w_N), \mathbf{v}(w_i)) \right)^{\frac{1}{N}}.$$

This model is inspired by results found in Section 3.2, which suggest that humans interpret whole-MWE composition as the geometric mean of the composition of its parts. While *maxsim* optimizes for higher scores of compositionality, and should thus yield better results for compositional MWEs, we hypothesize ($h_{\text{strat.geom}}$) that *geom* should obtain higher scores for idiomatic MWEs, due to its tendency to predict lower scores if either of the components is judged as non-compositional regarding the whole.

Other optimized functions such as the ones proposed by Yazdani, Farahmand, and Henderson (2015) could also be verified, but are out of the scope of this thesis, as they are based on supervised learning. A multiplicative version of *uniform* has also been considered in early experiments, but it did not present promising results, in particular in the case of sparse representations such as *PPMI-thresh*, in which the product of any two vectors tends to have an impractical number of non-zero dimensions.

The compositionality prediction model proposed in this thesis was implemented as part of the *mwetoolkit*⁴. The code is publicly available, and contains an internal module for the interpretation of multiple DSM formats, including sparse-context representations (such as the output of *minimantics*⁵) and dense representations (such as the output of *word2vec*⁶). Internally, the set of word vectors is represented as a sparse mapping from (target, context) word-form pairs to a real number; i.e. $(V, V) \rightarrow \mathbb{R}$. In the case of dense (fixed-length) input vectors, where there is no clear semantics attached to each dimension, we generate artificial identifiers for the context ($c_0, c_1, c_2 \dots c_{D-1}$). This allows a unified view of all types of vector representations.

In the implementation of the model, the sparse mapping from target-context pairs to a numerical representation is instantiated as a hash-table. Only non-zero mappings are explicitly represented, and thus all missing (target, context)

⁴ <http://mwetoolkit.sf.net>

⁵ <https://github.com/ceramisch/minimantics>

⁶ <https://code.google.com/archive/p/word2vec>

pairs are assumed to map to 0. When launching the module, a parameter can be specified to choose from among the aforementioned composition strategies. The compositionality scores are then predicted according to the model. Several types of output format can be specified, including a comprehensive XML format and a more lightweight CSV output. A full description of the compositionality prediction model and the associated tool has been published as Cordeiro, Ramisch, and Villavicencio (2016b).

4.3. Corpus preprocessing

Experiments in this thesis are based on distributional models built for English, French and Portuguese. The construction of these models uses the lemmatized and POS-tagged versions of the following corpora:

- For English, the ukWaC (Baroni, Bernardini, Ferraresi, et al. 2009), with 2.25 billion tokens, parsed with MaltParser (Nivre, Hall, and Nilsson 2006).
- For French, the frWaC, with 1.61 billion tokens preprocessed with TreeTagger (Schmid 1995).
- For Portuguese, a combination of brWaC (Boos, Prestes, and Villavicencio 2014), Corpus Brasileiro⁷ and all Wikipedia articles⁸, with a total of 1.91 billion tokens. This corpus was obtained in raw form, and parsed with PALAVRAS (Bick 2000) for the specific purpose of this thesis.

Target MWEs in these corpora are re-tokenized so as to be represented by a single token, with its components joined by an underscore character (e.g. the surface form EN *monkey business* → *monkey_business* and FR *belle-mère* → *belle_mère*).

During initial experiments, we noticed an inconsistency in the POS tags of MWE occurrences (e.g. the joined token *sitting_duck* had most of its occurrences tagged as VERB_NOUN instead of ADJ_NOUN). To handle such errors, we also re-tag every annotated occurrence of an MWE with a global manually selected POS tag.⁹

All forms are then lowercased (surface forms, lemmas and POS tags); and noisy tokens, with special characters, numbers or punctuation, are removed. Additionally, ligatures are normalized for French (e.g. *œ* → *oe*) and a spellchecker¹⁰ is applied to normalize words across English spelling variants (e.g. *color* →

⁷ <http://corpusbrasileiro.pucsp.br/cb/Inicial.html>

⁸ Wikipedia articles downloaded on June 2016.

⁹ For simplicity, our work assumes that every annotated occurrence is an instance of an MWE. We do not account for literal readings, such as cases of *sitting duck* that refer to an actual duck that is sitting.

¹⁰ <https://hunspell.github.io>

colour). Additionally, proper nouns are replaced by a placeholder to reduce data sparsity.

To evaluate the influence of preprocessing in model accuracy (see Section 5.3.1), we generated four versions of each corpus, with decreasing levels of specificity in the informational content of each token:

1. *surface*⁺: the surface-level forms of every word in the original corpus, with only the preprocessing described above. Example:

```
| she is not interested in your fake crocodile_tears !
```

2. *surface*: stopword removal¹¹; generating a corpus of surface forms for content words only (i.e. nouns, adjectives, adverbs and verbs). Example:

```
| is not interested fake crocodile_tears
```

3. *lemma_{POS}*: stopword removal, lemmatization¹² and POS tagging; generating a corpus of content words distinguished by POS tags, encoded in the format *lemma/tag*. This conflates the multiple inflectional forms of a word while maintaining the information of its grammatical category. Example:

```
| be/VERB not/ADV interested/ADJ fake/ADJ crocodile_tear/N_N
```

4. *lemma*: stopword removal and lemmatization without POS tagging; generating a corpus containing only lemmas of content words. This conflates identically-spelled words of different grammatical categories. Example:

```
| be not interested fake crocodile_tear
```

4.4. DSMs

One of the goals of this thesis is to evaluate the proposed model of compositionality prediction under a variety of distributional settings. In particular, we verify the impact of different types of DSMs (previously described in Section 2.4) in the predictive abilities of the model. For reproducibility, we present below the fixed parameters that were used in the DSM instantiations:¹³

PPMI We consider three DSMs based on positive pointwise mutual information (PPMI). In all cases, the representation of a target word is a vector containing the PPMI association scores between the target and its contexts. The contexts are nouns and verbs, selected in a symmetric sliding window of w words to the

¹¹ Stopword removal reduces the size of the corpus. Given that only nouns and verbs are used as contexts, the resulting co-occurrence matrices for *surface* will be less sparse than the matrices for *surface*⁺, for a given window size.

¹² In the lemmatized corpora, the lemmas of proper names are replaced by placeholders.

¹³ These parameters were selected with the goal of homogenizing the configurations across DSMs, and to follow the original paper’s recommendations in the cases where the default differs.

left/right and weighted linearly according to their distance d to the target (Levy, Goldberg, and Dagan 2015). We consider three models that differ in how the contexts are selected:

- *PPMI-thresh*, where the vectors are $|V|$ -dimensional but only the top δ local contexts with highest PPMI for each target word have non-zero values (Padró, Idiart, Villavicencio, et al. 2014).
- *PPMI-TopK*, where the vectors are k -dimensional, with a fixed global list of k words to be considered as context. We have defined k as the 1000 most frequent words in the corpus after removing the top 50 most frequent words, replicating the setup from Salehi, Cook, and Baldwin (2015).
- *PPMI-SVD*, where SVD is used to factorize the PPMI matrix and reduce its dimensionality from $|V|$ to D .¹⁴ We use a Context Distribution Smoothing of 0.75 and negative sampling of 5 for the SVD (Levy, Goldberg, and Dagan 2015).

w2v We perform experiments on both variants of word2vec (Mikolov, Sutskever, Chen, et al. 2013): continuous bag-of-words (*w2v-cbow*) and skip-gram (*w2v-sg*). The models are built with default configurations, except for the following: no hierarchical softmax; negative sampling of 25; frequent-word downsampling weight of 10^{-6} ; execution of 15 training iterations. We use the default minimum word count threshold of 5.

glove GloVe implements a factorization of the co-occurrence count matrix (Pennington, Socher, and Manning 2014). We use its default configurations, except for the following: internal cutoff parameter $x_{max} = 75$; co-occurrence matrix is built in 15 iterations. For lemma-based models, we use the minimum word count threshold of 5. Due to the large vocabulary size, we use a threshold¹⁵ of 15 for *surface* and 20 for *surface*⁺.

lexvec The lexvec model (Salle, Villavicencio, and Idiart 2016) factorizes the PPMI matrices, strongly penalizing prediction errors on frequent words. We use default configurations, except for the following: negative sampling of 25; subsampling threshold of 10^{-6} ; processes the corpus for 15 iterations. Due to the large vocabulary size, we use a minimum word count threshold of 10 for lemma-based models and 100 for *surface* and *surface*⁺.¹⁶

¹⁴ We use the hyperwords toolkit: <https://bitbucket.org/omerlevy/hyperwords>

¹⁵ Thresholds were selected so as to not use more than 128 GB of RAM during the construction of a DSM instance.

¹⁶ This is in line with the authors' threshold suggestion in their paper.

4.5. Parameters

For every DSM, we construct multiple distributional models under different sets of configurations. In particular, we exhaustively evaluate the influence of the following variables:

- **WORDFORM**: One of the four word-form and stopword removal variants when representing a corpus: *surface*⁺, *surface*, *lemma*, and *lemma*_{PoS} (see Section 4.3). These variants were selected so as to represent different levels of specificity in the informational content of the tokens.
- **WINDOWSIZE**: Indicates the number of context words that will be considered to the left/right of the target word when searching for target-context co-occurrence pairs. We evaluate the behavior of compositionality prediction when the underlying DSM model is built with context window sizes of 1+1, 4+4, and 8+8.¹⁷
- **DIMENSION**: We generate models with 250, 500 and 750 dimensions. The underlying hypothesis is that, the higher the number of dimensions, the more accurate the representation of the context is going to be.

Table 4.1 presents the set of all possible parameter configurations. These combinations produce a total of 228 models per language (12 models for *PPMI-TopK*, 36 models for each of the other 6 DSMs). Throughout the thesis, when referring to a particular model configuration, an abbreviated notation will be used. For example, **WORDFORM**=*lemma*, with **WINDOWSIZE**=4+4 and **DIMENSION**=250 will be represented as *lemma.w₄.d₂₅₀*.

DSM	DIMENSION	WORDFORM	WINDOWSIZE
<i>PPMI-TopK</i>	$D = 1000$	$surface^+,$ $surface,$ $lemma,$ $lemma_{PoS}$	1+1, 4+4, 8+8
<i>PPMI-thresh</i>	$D = V , \delta \in \{250, 500, 750\}$		
<i>PPMI-SVD</i>	$D \in \{250, 500, 750\}$		
<i>w2v-cbow</i>			
<i>w2v-sg</i>			
<i>glove</i>			
<i>lexvec</i>			

Table 4.1.: 228 parameter combinations across all DSMs.

¹⁷ Most works in the literature choose a window of size between the extremes 1+1 and 10+10, with a few works considering higher window sizes such as 16+16 or 20+20 (Kielar and Clark 2014; Laposa and Evert 2014).

4.6. Evaluation setup

For the intrinsic evaluation of the compositionality prediction model in Chapter 5, we calculate the compositionality scores for every MWE in a dataset and compare them to the human-rated scores. The following datasets are evaluated:

- For English: *Reddy*, *Reddy*⁺⁺, *EN-comp*₉₀ and *Farahmand*;
- For French: *FR-comp*;
- For Portuguese: *PT-comp*.

The datasets *Reddy* and *Farahmand* were described in Section 2.5.3. The other datasets were constructed as part of this thesis, and were described in Chapter 3.

For most datasets, we report Spearman’s ρ correlation between the ranking provided by humans and those calculated from the models (as explained in Section 2.2.5). Exceptionally for the *Farahmand* dataset, due to the binary nature of its compositionality scores, we follow Yazdani, Farahmand, and Henderson (2015) and report the best F₁ score (BF1), obtained by calculating the F₁ score for the top k MWEs classified as positive (non-compositional), for all possible values of k (described in Section 2.2.6).

Evaluation metrics were calculated for a total of more than 8 thousand models (see Figure 4.2). Given the high number of experiments performed, we report the best performance of each model parameter. For instance, the performances reported for *w2v-cbow* using different values of WINDOWSIZE are the best configurations across all possible values of other parameters (i.e. DIMENSION and WORDFORM). This avoids reporting local maxima that can arise if one fixes all other parameters when evaluating a given one (Lapesa and Evert 2014).




Figure 4.2.: Number of compositionality prediction models evaluated: $6 \times 228 \times 6 = 8\,208$.

For English datasets, we distinguish between *strict* and *fallback* evaluation. Strict evaluation corresponds to the performance of the model only on those MWEs that have a vector representation in all underlying DSMs: 89 (out of 90) for *Reddy*, 86 (out of 90) for *EN-comp*₉₀, 175 (out of 180) for *Reddy*⁺⁺, and 913 (out of 1042) for *Farahmand*. Fallback evaluation considers the full dataset, using a fallback strategy for the imputation of missing values, assigning the average of other compositionality scores to MWEs in which one of the vectors $\mathbf{v}(w_1)$, $\mathbf{v}(w_2)$ or \mathbf{v}_u has not been built due to a lack of occurrences in the corpus (Salehi, Cook, and Baldwin 2015). This distinction is particularly important in the case of *Farahmand*, which contains more rare MWEs¹⁸ such as *universe human* and *mankind instruction*, so that 129 MWEs do not occur often enough in the English corpus. Strict evaluation allows us to properly evaluate the quality of the predictive method itself, while fallback evaluation allows us to evaluate the quality of corpus + method, and is the better alternative when comparing to state-of-the-art results (as it considers the whole dataset). Only strict evaluation is reported for *FR-comp* and *PT-comp*, as all MWEs are frequent enough in their respective corpora.

To determine whether the results for different DSM configurations are statistically different from each other, we present results from Wilcoxon’s sign-rank test (Rey and Neuhäuser 2011).

¹⁸ Partly due to Wikipedia tokenization errors.

5. Intrinsic evaluation of compositionality prediction

This chapter presents an extensive intrinsic evaluation of the compositionality prediction framework presented in Chapter 4, using the three datasets whose construction was presented in Chapter 3. We construct DSM instances under multiple configurations for the three target languages. For each configuration, we generate compositionality predictions for all nominal compounds (NCs), which we then compare with the compositionality scores provided by humans.

This chapter is organized as follows: Section 5.1 presents our findings on the accuracy of compositionality prediction models using state-of-the-art DSMs for the representation of word semantics. Section 5.2 investigates the impact of DSM-specific parameters related to the size of the context window and the number of dimensions used to represent context. Section 5.3 examines the impact of corpus parameters related to corpus size and to the degree of corpus preprocessing adopted. Section 5.4 extends the evaluations performed on *uniform* prediction so as to encompass five other prediction strategies. Section 5.5 performs a variety of sanity checks involving other model-specific parameters. Section 5.6 presents an error analysis comparing predicted compositionality and different variables associated with the compounds. Finally, Section 5.7 summarizes the results from this chapter.

5.1. Overall highest results per DSM

This first evaluation aims at verifying whether some DSMs, independently of their specific parameters, are more suitable for a given dataset/language than others. We perform a language-based analysis, evaluating the highest-scoring parameter combination of each DSM. All evaluations reported here use the *uniform* composition strategy (described in Section 4.2).

5.1.1. English

Figure 5.1 presents the best scores achieved under each DSM for the English datasets. The predictions for the *Reddy*⁺⁺ dataset were evaluated through Spearman ρ , while *Farahmand* predictions were evaluated with BF_1 (both described in Section 2.2). Each of the wide bars in these graphs represents the highest score obtained from the set of 36 different configurations¹, using different combinations of WORDFORM, WINDOWSIZE, and DIMENSION. Similarly, each of the

¹ Only 12 configurations for *PPMI-TopK*, as the number of dimensions is fixed at 1000.

narrow inner bar represents the highest score among 36 configurations using fallback evaluation. While the fallback evaluation is responsible for slightly higher results in *Reddy*⁺⁺, its pessimistic approach is detrimental when evaluating the predictive model on the *Farahmand* dataset, which contains a considerable number of NCs that do not appear in the corpus. In both cases, these two types of evaluation produce similar rankings among the different DSMs, and we will henceforth focus on the highest results of strict evaluation (outer bars).

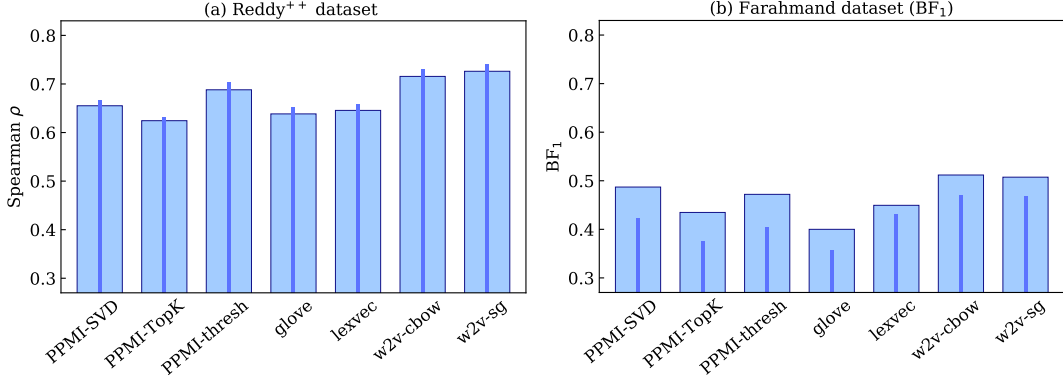


Figure 5.1.: Overall highest results per DSM on English datasets.

For the *Reddy*⁺⁺ dataset, the highest results are found for the word-embedding models of *w2v*: the highest-Spearman *w2v-sg* model has a $\rho = .741$, while the best *w2v-cbow* has a $\rho = .730$. These results are followed by *PPMI-thresh*, with $\rho = .704$. Other models offer progressively inferior results: *PPMI-SVD* ($\rho = .666$), *lexvec* ($\rho = .658$), *glove* ($\rho = .651$) and *PPMI-TopK* ($\rho = .632$). We performed Wilcoxon’s sign-rank test between all possible pairs of highest-Spearman configurations. The distributions of *w2v-cbow*, *w2v-sg* and *glove* were not deemed to be different from one another pairwise. Moreover, *glove* was not deemed different from *lexvec* or from *PPMI-SVD*. This is somewhat surprising, given the difference in scores between *glove* and the other DSMs (especially *w2v*). All other model pairs were deemed statistically different from each other ($p < 0.05$).

The *Reddy*⁺⁺ dataset combines all NCs from *Reddy* and *EN-comp*₉₀. If the *Reddy* dataset is considered by itself, we see the same trends as in *Reddy*⁺⁺. The overall best performance is found for *w2v-sg* ($\rho = .812$). The performance of the models *PPMI-thresh* ($\rho = .803$) and *w2v-cbow* ($\rho = .796$) closely follow the first place. If we isolate the *EN-comp*₉₀ NCs, a similar pattern emerges: the highest performance is achieved by *w2v-sg* ($\rho = .669$) and *w2v-cbow* ($\rho = .665$). Note that the highest results for *EN-comp*₉₀ are inferior to the ones obtained for *Reddy*. As we will see later, the best results for the French and Portuguese datasets are also in the same range as *EN-comp*₉₀. This difference in performance might be caused by the fact that these 3 datasets contain a higher amount of adjective+noun pairs than *Reddy*. As suggested in Section 3.2.2, humans had more difficulty judging the compositionality of adjectives than judging the compositionality of nouns.

This could imply that noun+adjective scores are less reliable than noun+noun scores, and thus automatic methods should also obtain lower scores when predicting the compositionality of adjectives.

Similarly to the *Reddy*⁺⁺ dataset, evaluation on *Farahmand* yields the best results for the models *w2v-sg* (strict $BF_1 = .498$, fallback $BF_1 = .455$) and *w2v-cbow* (strict $BF_1 = .501$, fallback $BF_1 = .471$). The highest results are comparable to the $BF_1 = .487$ reported by Yazdani, Farahmand, and Henderson (2015), while avoiding their use of functions whose parameters must be tuned through supervised learning for the prediction of compositionality.

5.1.2. French

As shown in Figure 5.2(a), overall *FR-comp* results in terms of Spearman correlation are reasonably different from *Reddy*⁺⁺ results. One of the most striking differences is the fact that the *w2v* models have lower quality. They are notably surpassed by *PPMI-thresh*, which rises to the first place with $\rho = .702$. This result is followed by word-embedding models: *glove* has $\rho = .680$ and *lexvec* has $\rho = .677$. Only then do we see the neural-network *w2v* models: *w2v-sg* ($\rho = .653$) and *w2v-cbow* ($\rho = .652$). Other PPMI-based models have a lower quality, dropping to the lowest $\rho = .550$ for *PPMI-TopK*.

As in the case of *Reddy*⁺⁺, we performed Wilcoxon’s sign-rank test between all model pairs ($p < 0.05$), and both *w2v* models were deemed equivalent (i.e. we could not reject the hypothesis that they followed the same distribution). The *glove* model, however, was deemed different from both *w2v* configurations. The *PPMI-SVD* model was deemed equivalent to all other models except for *lexvec*. All other model pairs were deemed different from each other. Particularly in the case of the highest-Spearman DSM, *PPMI-thresh*, these results confirm that its best configuration is responsible for compositionality predictions that are statistically different from the predictions of other models.




Figure 5.2.: Overall highest results per DSM on French and Portuguese datasets.

5.1.3. Portuguese

Figure 5.2(b) presents the overall highest Spearman correlations for the *PT-comp* dataset. As in the case of *FR-comp*, the *PPMI-thresh* model leads with the highest score ($\rho = .602$). This result is followed closely by word-embedding models: *w2v-cbow* has $\rho = .588$ and *w2v-sg* has $\rho = .586$. They are followed by *lexvec* with $\rho = .570$ and *glove* with $\rho = .555$. As with the French dataset, the other PPMI methods had the lowest scores: *PPMI-SVD* has $\rho = .530$ and *PPMI-TopK* has $\rho = .519$.

We performed Wilcoxon’s sign-rank test between all model pairs ($p < 0.05$). Just like for *Reddy++*, the distributions of *w2v-cbow*, *w2v-sg* and *glove* were deemed equivalent to each other pairwise. The *PPMI-SVD* and *glove* models were also deemed equivalent. Moreover, *PPMI-TopK* was deemed equivalent to all other models except *glove* and *PPMI-SVD*. Other model pairs were deemed different from each other ($p < 0.05$). As in the case of French results, *PPMI-thresh* had the highest scores, and Wilcoxon’s test has confirmed that its predictions are statistically different from the predictions of other models.

5.1.4. Cross-language analysis

Table 5.1 presents, for each dataset, the Spearman correlation score of the highest-Spearman configuration for every DSM (in strict/fallback format for English datasets), with the top strict score highlighted in bold. In most of the cases, the best score obtained under fallback evaluation is comparable to the best strict score. Fallback results are considerably lower in the case of *Farahmand*. This reflects the fact that this dataset is not balanced with regards to compositionality: most of its NCs are compositional, leading to a higher average compositionality score that may not be suitable for the missing NCs, as these tend to be idiomatic.

Dataset	<i>PPMI-SVD</i>	<i>PPMI-TopK</i>	<i>PPMI-thresh</i>	<i>glove</i>	<i>lexvec</i>	<i>w2v-cbow</i>	<i>w2v-sg</i>
<i>FR-comp</i>	.58	.55	.70	.68	.68	.65	.65
<i>PT-comp</i>	.53	.52	.60	.55	.57	.59	.59
<i>EN-comp</i> ₉₀	.59/.60	.56/.56	.59/.60	.52/.54	.55/.57	.65/.67	.65 /.67
<i>Reddy++</i>	.66/.67	.62/.63	.69/.70	.64/.65	.65/.66	.72/.73	.73 /.74
<i>Reddy</i>	.74/.74	.71/.72	.79/.80	.75/.76	.77/.77	.80/.80	.81 /.81
<i>Farahmand</i>	.49/.42	.43/.38	.47/.40	.40/.36	.45/.43	.51/.47	.51 /.47

Table 5.1.: Highest Spearman ρ for all datasets (strict/fallback format for English datasets). The *Farahmand* dataset uses the highest BF_1 instead.

Focusing on strict evaluation results, some interesting patterns can be observed in a comparison of the three languages. In all of the collected datasets, the predictions from *w2v-cbow* and *w2v-sg* follow the same distribution (as per a Wilcoxon sign-rank test), which is reflected in the fact that their scores are almost identical.

Concerning the highest Spearman score out of all DSMs, the *w2v* models yield the best results for the two English datasets, while the highest scores for the two Romance languages were attained instead by *PPMI-thresh*. On the other side of this scale, the *PPMI-TopK* model is consistently ranked among the worst results. The best *PPMI-SVD* configuration presents a similar behavior, with consistently low results for all datasets but *Farahmand*.

5.2. DSM parameters

In this section, we investigate the hypothesis $h_{\text{accur}} \leftarrow \text{DSM}$, which affirms that the accuracy of the model depends on DSM-specific parameters. We consider two parameters that can be independently tuned in every DSM: context-window size and number of dimensions in the output vectors.

5.2.1. Context-window size

DSMs build the representation of every word based on the frequency of other words that appear in its context. A very simple way of defining such a context is through a window, whereby the context of a word with e.g. `WINDOWSIZE=4+4` would consist in the previous four words and the following four words in the text. Most works in the literature construct DSMs with window sizes between the extremes `1+1` and `10+10`, with a few works considering larger window sizes such as `16+16` or `20+20` (Kiela and Clark 2014; Lapesa and Evert 2014). We evaluate the behavior of compositionality prediction when the underlying DSM model is built with the commonly-used context-window sizes of `1+1`, `4+4`, and `8+8`.² Our hypothesis ($h_{\text{accur}} \leftarrow \text{DSM.window}$) is that the highest scores should be obtained by window sizes of `8+8`, as the extra amount of data would lead to a better representation of the word-level semantics.

As can be seen in Figure 5.3, the performance with different context-window sizes is mostly DSM-dependent. In the case of *PPMI-SVD*, a window of size `1+1` yields better results for all datasets, with the exception of *Farahmand* evaluations. The *glove* model exhibits the opposite behavior: windows of size `1+1` are consistently worse than the windows of size `4+4` or `8+8`. These results seem to be related to the manner with which the weights decay for different models. In the case of *PPMI-SVD* with `WINDOWSIZE=8+8`, a context word at distance d from its target word is weighted as $\frac{8-d}{8}$. In the case of *glove*, the decay happens much faster, with a weight of $\frac{8}{d}$, which allows the model to look farther away without being affected by the extra noise associated with the more distant contexts. For the other DSMs, window size is not a visible predictor of performance. In the exceptional case of *PPMI-thresh*, the results were language-dependent instead: French and Portuguese data can be better approximated through smaller

² Section 5.5.3 also performs some sanity checks for windows of size `2+2`.




Figure 5.3.: Best Spearman's ρ per DSM and WindowSize.

windows, while English data (including *Reddy* and *Reddy*⁺⁺) displays a weaker preference for larger window sizes. The appropriate choice of window size has been shown to be task-specific (Lapesa and Evert 2017), and the results above suggest that, in the task of compositionality prediction, even the choice of DSM may interact with this parameter.

5.2.2. Number of dimensions

When instantiating a DSM, there is a trade-off in the number of vector dimensions. Models that have lower amount of dimensions will correspondingly have a smaller memory footprint³, while models that have a larger number of dimensions eschew any memory concerns so as to be able to represent more fine-grained patterns of co-occurrence. The question is whether these extra dimensions can be put to good use by state-of-the-art DSMs. Most works in the literature build distributional models whose vectors contain between 200 and 900 dimensions (Baroni, Dinu, and Kruszewski 2014; Lapesa and Evert 2014). We evaluate different DSMs by using DIMENSION=250, which approximates the common value used in the literature. We additionally present results for two

³ Memory usage grows linearly with the number of dimensions.

of its multiples: 500 and 750 dimensions. Our hypothesis ($h_{\text{accur}} \leftarrow \text{DSM.dims}$) is that the highest scores should be obtained by DSMs where the vectors contain a higher number of dimensions.

As Figure 5.4 shows, for most DSMs, an increase in the number of dimensions causes a moderate increase in the quality of the predictive model, reflecting the additional information that can be used to perform the compositionality predictions. This is particularly true in the case of the DSMs that obtain the highest Spearman scores. This behavior is however inverted for *PPMI-SVD*, in which the highest results can be obtained by building lower-dimension models. Moreover, two DSMs seem to be particularly unaffected by the number of dimensions: *glove* and *lexvec* models seem to have around the same predictive power regardless of the number of dimensions. Overall, these results suggest that, across all DSMs, the best scores can be obtained by configurations involving a higher number of dimensions, as hypothesized.




Figure 5.4.: Best Spearman's ρ per DSM and Dimension.

5.3. Corpus parameters

According to the hypothesis $h_{\text{accur}} \leftarrow \text{corpus}$, the accuracy of the semantic representation in a DSM is dependent on the quality of the input representation. In this

section, we analyze the impact of different types of corpus preprocessing, corpus size variation, as well as the use of parallel sub-corpora in the prediction of NC compositionality.

5.3.1. Type of preprocessing

Most works in the literature on distributional models focus on the analysis and tuning of different statistical parameters, reducing the preprocessing to tokenization, along with simple procedures such as lower-casing and rare word removal (Mikolov, Sutskever, Chen, et al. 2013; Pennington, Socher, and Manning 2014; Levy, Goldberg, and Dagan 2015; Salle, Villavicencio, and Idiart 2016). Works that rely on DSMs in semantic tasks tend to consider other preprocessing techniques, such as lemmatization, stemming, POS tagging, and stopword removal (Bulinaria and Levy 2012; Kiela and Clark 2014). The goal of such procedures is to increase the quality of the word representations, by conflating different uses of the same word into a single canonical form, by allowing the disambiguation of homonyms based on their syntactic function, and through the elimination of random noise from the corpora.

As described in Section 4.3, we consider four different levels of corpus preprocessing: $\text{WORDFORM}=\text{surface}^+$, *surface*, $\text{lemma}_{\text{POS}}$ and *lemma*. Under each of these configurations, there is a difference in how much information is condensed into each token in the corpus, with surface^+ being the most specific (every word in the corpus is considered verbatim) and *lemma* being the most general (where only the lemmas of content words are used in the representation of each token). Our hypothesis ($h_{\text{accur}} \leftarrow \text{corpus.wordform}$) is that the less specific configurations present a less sparse view of the data, contributing to higher-quality DSM representations and thus achieving higher scores.

Figure 5.5 presents the impact of different types of corpus preprocessing on the quality of the compositionality prediction model. The results seem to be language-dependent: The results for the English-language datasets are quite heterogeneous, while for the other two languages, the lemma-based word representations consistently allow a better prediction of compositionality scores. This phenomenon may be explained by the fact that French and Portuguese are morphologically richer than English. For the former languages, lemma-based representations reduce the sparsity in the data and allow more information to be gathered from the same amount of data. In the case of English, lemmatization has a reduced effect, and in particular for *PPMI-SVD*, it visibly reduces the quality of predictions.

The results for $\text{lemma}_{\text{POS}}$ represent the addition of POS tags to every word in the corpus. This extra information does not show any improvement over *lemma*. This suggests that words that share the same lemma are semantically close enough that any gains from disambiguation are compensated by the sparsity of a higher vocabulary size.




Figure 5.5.: Best Spearman's ρ per DSM and WordForm

The results obtained with *surface*⁺ are surprisingly similar to *surface*, which confirms previous suggestions that stopwords removal does not significantly affect the data (Bullinaria and Levy 2012). These results are achieved even though *surface*⁺ contains stopwords, which one might expect would dilute the DSM representation (due to their low level of association with most other words), which might then reduce the accuracy of predictions. A possible explanation could be that the stopwords in *surface*⁺ effectively contribute to a reduced window size of content words⁴, which is shown in Section 5.2.1 to consistently yield better results. Indeed, when looking at the highest *surface*⁺ scores across all models and datasets, the majority of the configurations involve WINDOWSIZE=1+1, further highlighting the role of the context-window size in the best-performing models.

5.3.2. Corpus size

In this thesis, we have presented results for compositionality prediction based on three similarly-sized corpora. The general intuition gathered from the literature is that predictions based on larger corpora should obtain higher scores, as

⁴ This could be investigated in future work with a WORDFORM that only uses the content words inside a content+stopword window.

rarer contexts would also be taken into account, thus improving the quality of the vector representation of these words. We express this intuition through the hypothesis $h_{\text{accur}} \leftarrow \text{corpus.size}$, which predicts higher scores for DSMs instantiated for larger-sized corpora.

This section performs a quantitative analysis of the impact of different corpus sizes on the quality of the predictions. For each of the *Reddy*, *FR-comp* and *PT-comp* datasets, we consider the highest-Spearman *PPMI-thresh* and *w2v-sg* configuration obtained thus far for the full-size corpus.⁵ We then build new models under the same configuration, but using corpus fragments of size varying from 1% to 100% of the whole corpus, increasing by steps of 1/100 at a time.

Figure 5.6(a) presents three graphs of the scores obtained by building models for the best *PPMI-thresh* configuration of each dataset. The 100 positions in the x-axis correspond to the corpus sizes (1% to 100%). Eight different samplings of corpus fragments were performed (for a total of 800 models per language), with each y-axis data point presenting the average of the 8 Spearman scores obtained from those samplings. Each data point also presents the sample standard deviation for those 8 executions. Points to the left of the vertical bar have at least one sampling with missing compounds, while points to the right have 100% of the compounds in all 8 samplings. The results suggest that, for the three languages, a corpus size of around 800 million to 1 billion tokens (40% of the whole corpus size) is large enough to obtain the best results, with further increases in the amount of data available only contributing marginally to the overall quality of the predictions.

Figure 5.6(b) presents three graphs with the scores obtained by the best *w2v-sg* configuration for each dataset. Due to the fact that *w2v-sg* is much more time-consuming than *PPMI-thresh*, a single sampling was used, and thus only one execution was performed for each datapoint (for a total of 100 vector models per language). Similarly to *PPMI-thresh*, a corpus fragment of around 40% size (800 million to 1 billion tokens) was already large enough for the results to stabilize close to the score obtained by the fragment of size 100%. This suggests that corpus size does strongly affect the quality of the underlying DSM representation, but that it reaches a plateau around a billion tokens. Even higher corpus sizes would presumably only offer a minor improvement in DSM representation quality. Future work should investigate whether a similar plateau can also be observed for other methods of compositionality prediction.

5.3.3. Parallel predictions

One idea that has been employed in the literature is that of *ensemble methods*, in which the predictions from multiple methods (e.g. multiple DSM instances built

⁵ Results for corpus size do not consider variations in the configurations, as the best configuration for the full-size corpus is used for every corpus size.

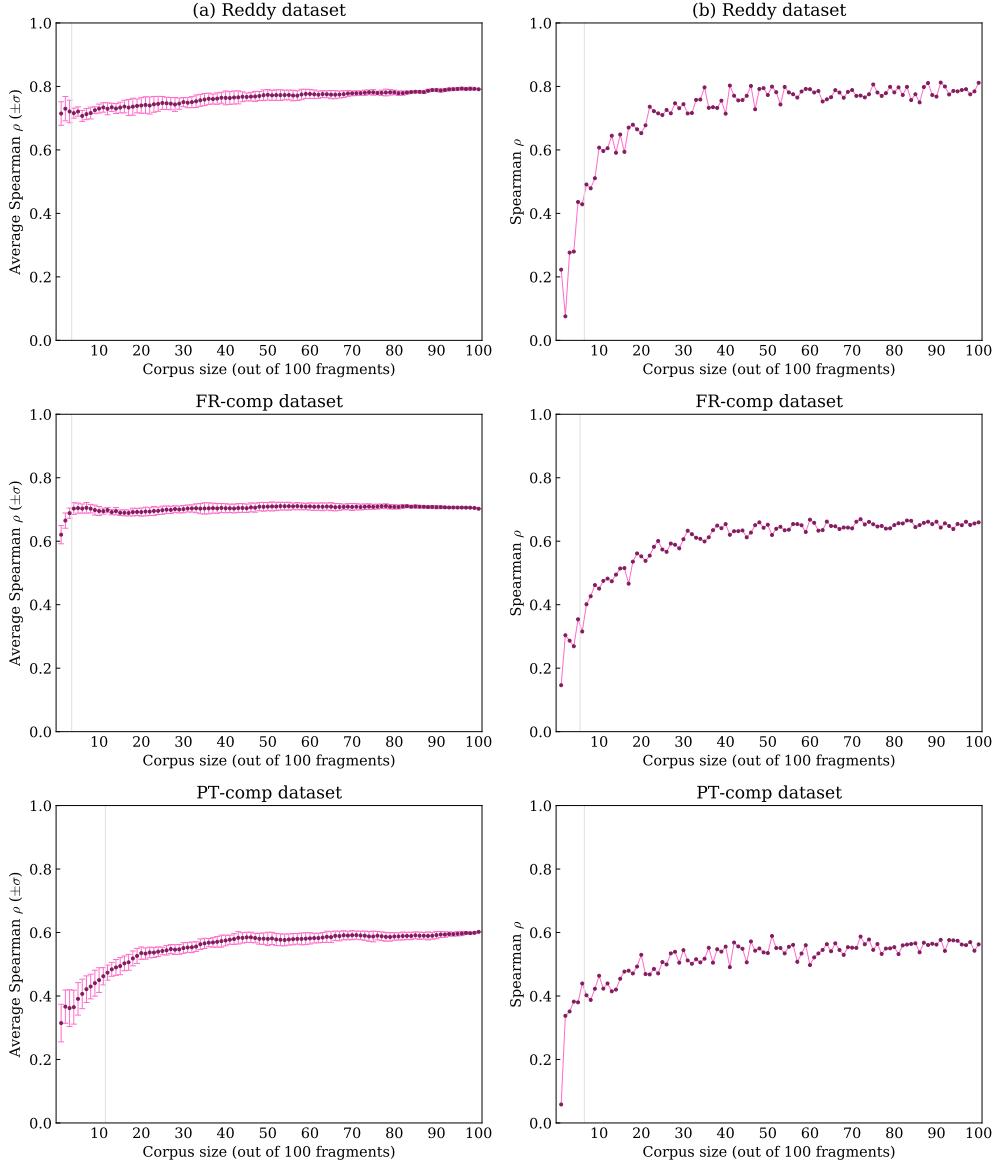


Figure 5.6.: Spearman’s ρ for different corpus sizes, running *PPMI-thresh* (left) and *w2v-sg* (right).

from the same data) are combined into a single prediction that outperforms all other methods (Zhou 2012). We leave ensemble experiments for future work, but we focus on an approach that is inspired by its success: *parallel predictions*.

The results obtained for different corpus sizes above imply that a subset of the corpus can lead to results that are equivalent to the ones obtained for the whole corpus. In fact, even when a smaller fraction of the corpus is considered, such as 20% of the total corpus size, the resulting model can still yield reasonably good predictions of compositionality. We hypothesize ($h_{\text{accur}} \leftarrow \text{corpus.parallel}$) that, just as an ensemble of *methods* instantiated from a single corpus may complement

each other and achieve higher scores, so can a single DSM method instantiated multiple times in parallel from an ensemble of *corpora* be combined so as to achieve accurate DSM representations.

We thus propose a technique in which the whole corpus is divided in M parallel fragments ($c_1, c_2, \dots c_M$). We then instantiate the same DSM M times in parallel, each one based on a different corpus fragment. These DSMs can then be used to produce a set of M compositionality predictions per compound, using the *uniform* strategy ($CS_{\beta(1)}, CS_{\beta(2)}, \dots CS_{\beta(M)}$). The M parallel predictions for each compound are then combined through the arithmetic average into a single compositionality score CS_P . We can then evaluate the set of CS_P for every compound by comparing them with the reference dataset.

In order to verify the hypothesis that the parallel predictions can yield results that are comparable to the ones obtained on the whole corpus, we ran an experiment in which the whole corpus was divided in $M = 5$ fragments of equal size, each corresponding to 20% of the whole corpus. Table 5.2 presents the Spearman scores of whole-corpus ($\rho_{100\%}$) and parallel prediction ($\rho_{5 \times 20\%}$) for two models: *PPMI-thresh* and *w2v-sg*. For the latter, we have considered smaller subsampling sizes rates of 10^{-3} and 10^{-4} , to account for the smaller corpus sizes. The results suggest that the parallel prediction on smaller corpus fragments can be as effective as a single prediction generated from the whole corpus.

Model (<i>Reddy</i> ⁺⁺)	$\rho_{100\%}$	$\rho_{5 \times 20\%}$	Difference (%)	Worst $\rho_{20\%}$	Best $\rho_{20\%}$
<i>PPMI-thresh</i>	.699	.680	(−1.9)	.626	.678
<i>w2v-sg</i> (10^{-3})	.731	.719	(−1.2)	.668	.708
<i>w2v-sg</i> (10^{-4})	.731	.717	(−1.4)	.667	.707
Model (<i>FR-comp</i>)	$\rho_{100\%}$	$\rho_{5 \times 20\%}$	Difference (%)	Worst $\rho_{20\%}$	Best $\rho_{20\%}$
<i>PPMI-thresh</i>	.702	.709	(+0.7)	.686	.714
<i>w2v-sg</i> (10^{-3})	.672	.685	(+1.3)	.654	.688
<i>w2v-sg</i> (10^{-4})	.672	.688	(+1.6)	.671	.693
Model (<i>PT-comp</i>)	$\rho_{100\%}$	$\rho_{5 \times 20\%}$	Difference (%)	Worst $\rho_{20\%}$	Best $\rho_{20\%}$
<i>PPMI-thresh</i>	.602	.572	(−3.1)	.496	.549
<i>w2v-sg</i> (10^{-3})	.586	.581	(−0.5)	.528	.569
<i>w2v-sg</i> (10^{-4})	.586	.581	(−0.6)	.520	.566

Table 5.2.: Results for whole-corpus and parallel predictions.

Table 5.2 also indicates the Spearman scores obtained for the highest and lowest-ranking corpus fragments (Worst $\rho_{20\%}$ and Best $\rho_{20\%}$, each using only 20% of the corpus). In the case of *Reddy*⁺⁺ and *PT-comp*, in all configurations, we can see that the average score obtained from the 5 fragments ($\rho_{5 \times 20\%}$) is slightly *higher* than the score obtained by the best fragment. This suggests that the technique of parallel predictions is actually able to combine the results from the different fragments, confirming the underlying hypothesis. We leave it for future

work the investigation of different ways of combining the parallel predictions into a single score.

One of the greatest advantages of parallel prediction is its potential for scalability. While a standard DSM-based predictive model requires the whole corpus to be processed at once, a parallel model allows the computation of such predictions in a distributed fashion. This reduces the total execution time, allows the better utilization of distributed resources (such as computer clusters), and bypasses memory limitations of a single machine.

5.4. Prediction strategy

Now that we have evaluated the impact of DSM and corpus parameters on the predicted compositionality scores, we turn to the underlying prediction strategy itself. As we have seen in Chapter 3, the elements of an MWE may vary in terms of the semantic contribution of each element to the MWE as a whole, and this may have an impact on the success of the composition model adopted for deriving the vector space representation of the MWE (hypothesis h_{strat}). For instance, adopting a uniform (50%:50%) composition for the elements of a compound might not accurately capture a faithful representation of compounds whose meaning is more semantically related to one of the components than to the other (as in the case of the compound *crocodile tears*, regarding its head, and *night owl*, regarding its modifier).

We compare below six different compositionality prediction strategies (all described in Section 4.2). Some of these strategies consider different variations of weights on the compound elements themselves: *uniform* uses a 50%:50% scheme, while two other strategies (*head* and *mod*) use a 0%:100% scheme. We also evaluate a proposed new model of additive composition, *maxsim*, which dynamically determines weights so as to assign an optimal proximity of the compound to each of its single-word elements. Additionally, we consider the *arith* and *geom* prediction strategies, in which the representation of the compound is independently compared to the representation of each component, and with the resulting score being the (arithmetic or geometric) mean of the comparison scores.

Table 5.3 presents the scores obtained for all strategies on the configurations in which *uniform* obtains its highest scores (i.e. using the best configuration for each DSM as reported up to now). In most of the cases, the score obtained for the *uniform* prediction is higher than both the *head* and *mod* scores when taken separately, which is in line with the hypothesis $h_{\text{strat}, \text{partial-info}}$ that these two strategies are somewhat limited due to the fact that they only consider half of the available distributional information.⁶ However, this difference is only moderate, with *mod*

⁶ Future work should investigate why the BF_1 scores of *Farahmand* are higher for *mod* than for all other strategies.

predictions in particular attaining results that are quite close to *uniform*.

Results for other strategies are slightly worse than *uniform*, suggesting that these approaches are either subpar, or that they improve in configurations that differ from the best *uniform* configurations.

Dataset	DSM	configuration	<i>uniform</i>	<i>maxsim</i>	<i>geom</i>	<i>arith</i>	<i>head</i>	<i>mod</i>
<i>Reddy</i>	<i>w2v-sg</i>	<i>surface.w₁.d₇₅₀</i>	.812	.802	.756	.805	.635	.752
<i>EN-comp₉₀</i>	<i>w2v-cbow</i>	<i>lemma.w₄.d₅₀₀</i>	.653	.651	.600	.647	.463	.613
<i>Reddy⁺⁺</i>	<i>w2v-sg</i>	<i>surface⁺.w₁.d₇₅₀</i>	.726	.730	.657	.718	.524	.677
<i>FR-comp</i>	<i>PPMI-thresh</i>	<i>lemma_{POS}.w₁.d₇₅₀</i>	.702	.688	.668	.698	.605	.603
<i>PT-comp</i>	<i>PPMI-thresh</i>	<i>lemma_{POS}.w₁.d₇₅₀</i>	.602	.577	.524	.595	.524	.413
<i>Farahmand</i>	<i>w2v-cbow</i>	<i>lemma_{POS}.w₈.d₂₅₀</i>	.501	.484	.523	.517	.402	.534

Table 5.3.: Spearman scores for best *uniform* model of each dataset, using different prediction strategies. The *Farahmand* dataset uses BF₁.

To verify whether the other prediction strategies improve models that differ from the highest-Spearman *uniform* configurations, we have evaluated every strategy on all DSM instances. Table 5.4 presents a summary of the highest scores obtainable for each prediction strategy individually (i.e. each score represents the best configuration for a given strategy evaluated on a given dataset). Every best score is statistically different from all other scores in its row ($p < 0.05$). Similarly to the above results, the score obtained for the *uniform* prediction is higher than the one obtained for both *head* and *mod* strategies (hypothesis $h_{\text{strat.partial-info}}$), which further suggests that the quality of *uniform* predictions is derived from the combination of the vector representation of the two words (in particular from *mod*).

The *arith* strategy obtains performance results that are very similar to the ones of *uniform*, reflecting the fact that both methods rely on an additive model of composition. Indeed, if we consider the average (across the 7 DSMs) of the Pearson correlation between the 180 predictions for the highest-Spearman configuration of *arith* and *uniform*, we obtain $r = .972$ for *Reddy⁺⁺*, $.991$ for *FR-comp* and $.969$ for *PT-comp*, confirming that these models produce very similar predictions. Moreover, these two strategies behave very similarly when we consider the Spearman scores obtained for the 228 DSM instances in each language: if we consider the Pearson correlation between the 228 pairs of Spearman scores associated with both strategies, we obtain $r = .981$ for *Reddy⁺⁺*, $.985$ for *FR-comp* and $.944$ *PT-comp*.

When one considers the highest-Spearman configuration for each strategy, results for *maxsim* are competitive with the results for *uniform*. While *maxsim* fares slightly better on English continuous-score datasets, *uniform* obtains slightly higher scores on the other two languages. Implicit to the calculation of *maxsim* is the assignment of weights for the components of every NC, and we have considered whether the assigned weights actually differ from the 50:50 assignment

<i>Dataset</i>	<i>uniform</i>	<i>maxsim</i>	<i>geom</i>	<i>arith</i>	<i>head</i>	<i>mod</i>
<i>Reddy</i>	.812	.814	.797	.805	.654	.776
<i>EN-comp₉₀</i>	.653	.659	.600	.647	.483	.615
<i>Reddy⁺⁺</i>	.726	.730	.677	.718	.555	.677
<i>FR-comp</i>	.702	.693	.699	.703	.617	.645
<i>PT-comp</i>	.602	.590	.580	.598	.558	.486
<i>Farahmand</i>	.501	.487	.529	.518	.422	.528

Table 5.4.: Highest Spearman score for each prediction strategy individually. *Farahmand* uses BF_1 .

of *uniform*. This does seem to be the case, as some NCs have a weight prediction that is much closer to human head-only and modifier-only scores than a 50:50 prediction. For example, in the *w2v-sg* prediction for *Reddy⁺⁺*, some NCs were weighted more heavily in favor of the head (e.g. *silver screen* had weight 11:89), while others had more weight in the modifier (e.g. *spelling bee* with weight 94:06), and with many intermediary NCs in between (e.g. *dirty word* with weight 47:53).

One of the side effects of calculating these weights is that they also reveal any bias in the semantic influence of the head or the modifier of each compound, and we consider whether this bias may be affecting the results on each dataset. Table 5.5 presents the highest-Spearman *maxsim* model for each dataset, along with the average of the weights assigned to head and modifier for every NC in the dataset. The results are extremely stable: while the weights that optimize for compositionality are fairly similar for the English datasets, they are highly discrepant for both *FR-comp* and *PT-comp*, in which the weight of the head is disproportionately higher than the weight of the modifier.

The fact that *FR-comp* and *PT-comp* compounds have the potential for higher compositional interpretation in the head than in the modifier could be elucidated by the consideration that all of the modifiers in these datasets are adjectives, while English-language modifiers may also be nouns. Therefore, the contribution of adjectives to the overall meaning could be lower due to some linguistic phenomenon. For example, some of the adjectives used in these compounds are highly polysemous, and could be seen contributing to some specific meaning is not found on isolated occurrences of the adjective itself (e.g. FR *beau* (lit. *beautiful*) is used in the translation of most *in-law* family members, such as *beau-frère* ‘brother-in-law’ (lit. *beautiful-brother*)).

Up to this point, it is still unclear whether it is true that the *maxsim* strategy is able to more aptly capture the semantics of compositional MWEs (hypothesis $h_{\text{strat.maxsim}}$). In order to better understand the behavior of *maxsim* with regards to *uniform*, we rank the compounds in *ALL-comp* according to three possible sets of scores: (a) the compositionality score assigned by human annotators; (b) the highest-Spearman *maxsim* prediction; and (c) the highest-Spearman *uniform*

Dataset	DSM	configuration	maxsim	weight _{head}	weight _{mod}
<i>Reddy</i>	<i>w2v-sg</i>	<i>surface</i> ⁺ .w ₁ .d ₇₅₀	.814	53	47
<i>EN-comp₉₀</i>	<i>w2v-cbow</i>	<i>lemma</i> .w ₈ .d ₇₅₀	.659	54	46
<i>Reddy</i> ⁺⁺	<i>w2v-sg</i>	<i>surface</i> ⁺ .w ₁ .d ₇₅₀	.730	55	45
<i>FR-comp</i>	<i>PPMI-thresh</i>	<i>lemma</i> .w ₁ .d ₇₅₀	.693	68	32
<i>PT-comp</i>	<i>w2v-sg</i>	<i>lemma</i> .w ₈ .d ₇₅₀	.590	68	32

Table 5.5.: Average weight of highest-Spearman *maxsim* model for each dataset.

prediction. Each compound is then assigned three corresponding ranks (positive integers): rk_{human} , rk_{maxsim} , $rk_{uniform}$. We then calculate the improvement score of each compound as:

$$\text{improv}_{maxsim} = |rk_{uniform} - rk_{human}| - |rk_{maxsim} - rk_{human}|.$$

Figure 5.7 presents the distribution of rank improvement scores for the highest-scoring *PPMI-thresh* and *w2v-sg* configurations.⁷ Each graph presents the improvement score for NCs from the three languages, ranked according to rk_{human} . It can be seen that, for most NCs, there is only a light variation in the rank compositionality of *maxsim*. For the NCs that have a more drastic variation in rank, positive improvements are associated with higher human-ranked compositionality (right side of the graph), while negative improvement scores are associated with idiomatic NCs. This confirms the hypothesis that *maxsim* can better capture the semantics of compositional MWEs, albeit this only applies to some outlier cases.




Figure 5.7.: Distribution of improv_{maxsim} as a function of human judgments.

Figure 5.8 presents the distribution of rank improvements for all NCs, ranked according to $rk_{uniform}$ instead. Differently from above, NCs with the highest

⁷ We focus on one representative of PPMI-based DSMs and one representative of word-embedding ones. Similar results were observed for the highest-Spearman configuration of other DSMs.

variation in rank are found on the left side of the graph, indicating that they were all initially judged as idiomatic. This indicates that *maxsim* tends to improve the score of NCs that humans considered more compositional, but that the *uniform* system considered more idiomatic. On the other hand, NCs that are correctly classified as idiomatic by the *uniform* prediction are somewhat under-estimated by *maxsim*. The positive and negative improvements are somewhat balanced, which explains why *maxsim* predictions fare as well as *uniform*.




Figure 5.8.: Distribution of $\text{improv}_{\text{maxsim}}$ as a function of *uniform* scores.

Figures 5.7 and 5.8 also indicate the outlier NCs with highest most improvement (numbers from 1 to 8), as well as the NCs with lowest improvement scores (letters from A to H). Table 5.6 presents these outlier NCs along with their improvement scores (see the Appendices A, B and C for the translation, glosses and human scores associated with these compounds). We can see that there is a disproportionate amount of outlier NCs for Portuguese and French (particularly the former), suggesting that *maxsim* has a stronger impact on those languages than on English. It is also noticeable that some NCs had a similar improvement score under both DSMs, with e.g. high improvement for PT *caixa forte* and low improvement scores for PT *coração partido*. It is further remarkable that equivalent NCs in different languages are similarly impacted by *maxsim*, as in the case of PT *caixa forte* and FR *coffre fort*. Nevertheless, *maxsim* does not present a considerable overall impact on the rank of the predictions, obtaining an average improvement of $\overline{\text{improv}_{\text{maxsim}}} = +0.41$.

As in the case of *maxsim*, we also consider the rank improvement of *geom* predictions over *uniform*. We rank the compounds in *ALL-comp* according to three possible sets of scores: (a) the compositionality score assigned by human annotators; (b) the highest-Spearman *geom* prediction; and (c) the highest-Spearman *uniform* prediction. Each compound is then assigned three corresponding ranks (positive integers): rk_{human} , rk_{geom} , $\text{rk}_{\text{uniform}}$. We then calculate the improvement score of each compound as:

$$\text{improv}_{\text{geom}} = |\text{rk}_{\text{uniform}} - \text{rk}_{\text{human}}| - |\text{rk}_{\text{geom}} - \text{rk}_{\text{human}}|.$$

ID	improv	PPMI-thresh	improv	w2v-sg
1	(+90)	FR <i>premier plan</i>	(+138)	PT <i>cerca viva</i>
2	(+88)	FR <i>matière première</i>	(+126)	FR <i>coffre fort</i>
3	(+86)	PT <i>amigo oculto</i>	(+116)	PT <i>caixa forte</i>
4	(+67)	FR <i>première dame</i>	(+107)	PT <i>golpe baixo</i>
5	(+63)	PT <i>caixa forte</i>	(+100)	PT <i>primeira necessidade</i>
6	(+58)	PT <i>prato feito</i>	(+95)	EN <i>role model</i>
7	(+53)	FR <i>idée reçue</i>	(+79)	FR <i>bonne pratique</i>
8	(+48)	FR <i>marée noire</i>	(+69)	PT <i>carta aberta</i>
H	(-42)	PT <i>alta costura</i>	(-68)	FR <i>bras droit</i>
G	(-44)	EN <i>half sister</i>	(-70)	PT <i>alta costura</i>
F	(-44)	EN <i>melting pot</i>	(-71)	PT <i>carne vermelha</i>
E	(-46)	FR <i>berger allemand</i>	(-82)	PT <i>alto mar</i>
D	(-52)	PT <i>mar aberto</i>	(-85)	PT <i>mesa redonda</i>
C	(-55)	PT <i>febre amarela</i>	(-86)	EN <i>half sister</i>
B	(-81)	PT <i>livro aberto</i>	(-109)	PT <i>febre amarela</i>
A	(-83)	PT <i>coração partido</i>	(-128)	PT <i>coração partido</i>

Table 5.6.: Outliers regarding positive/negative *maxsim* improvement.

The hypothesis ($h_{\text{strat.geom}}$) is that *geom* should more accurately represent the semantics of idiomatic NCs, and this would be reflected in the improvement scores. Figure 5.9 presents the distribution of rank improvements for NCs in *ALL-comp*, ranked according to rk_{human} . It can be seen that, for most NCs, there is only a slight variation in the rank compositionality of *geom*. For the NCs that have a more drastic variation in rank, positive improvements are slightly more associated with lower human-ranked compositionality (left side of the graph), while negative improvement scores are visibly associated with compositional NCs (right side of the graph). This is the opposite of what was observed for *maxsim*, and confirms the interpretation that these models optimize for opposite extremes of compositionality (with *geom* focusing on idiomatic NCs at the expense of more compositional ones). As in the case of *maxsim*, this behavior is only observed for the outlier cases.

Figure 5.10 presents the distribution of rank improvements for all NCs in the highest-Spearman configuration, ranked according to $\text{rk}_{\text{uniform}}$ instead. Here again, the behavior of the *geom* strategy is the opposite of what was observed for *maxsim*: NCs with the highest variation in rank are found on the right side of the graph, indicating that they were all initially judged as compositional. This indicates that *geom* tends to improve the score of NCs that humans considered more idiomatic, but that the *uniform* system considered more compositional. On the other hand, NCs that are correctly classified as compositional by the *uniform* prediction are somewhat pessimized by *geom*.

Figures 5.9 and 5.10 also indicate the outlier NCs with the highest and lowest improvement scores (numbers and letters, respectively). Table 5.7 presents these




Figure 5.9.: Distribution of $\text{improv}_{\text{geom}}$ as a function of human judgments.




Figure 5.10.: Distribution of $\text{improv}_{\text{geom}}$ as a function of *uniform* scores.

outlier NCs along with their improvement scores. As in the case of *maxsim*, the majority of the outliers belong to the Portuguese dataset. Some of the NCs that were found as outliers in *maxsim* re-appear as outliers for *geom* with inverted polarity in the improvement score, e.g. FR *bras droit* scores predicted by *PPMI-thresh* ($\text{improv}_{\text{maxsim}} = +58$, $\text{improv}_{\text{geom}} = -234$) and PT *prato feito* as predicted by *w2v-sg* ($\text{improv}_{\text{maxsim}} = -68$, $\text{improv}_{\text{geom}} = +228$). This suggests that future work should consider combining both approaches into a single prediction strategy that decides which sub-strategy to use as a function of the *uniform* prediction for each NC. As it stands, however, the *geom* strategy has a mild negative influence on the rank of the predictions, obtaining an average improvement score of $\overline{\text{improv}_{\text{geom}}} = -7.87$.

ID	improv	<i>PPMI-thresh</i>	improv	<i>w2v-sg</i>
1	(+157)	EN <i>snail mail</i>	(+228)	FR <i>bras droit</i>
2	(+110)	FR <i>guerre civile</i>	(+158)	PT <i>lua nova</i>
3	(+109)	FR <i>disque dur</i>	(+127)	PT <i>alto mar</i>
4	(+104)	PT <i>alto mar</i>	(+104)	PT <i>pé direito</i>
5	(+93)	PT <i>ônibus executivo</i>	(+89)	EN <i>carpet bombing</i>
6	(+85)	EN <i>search engine</i>	(+75)	PT <i>lista negra</i>
7	(+82)	PT <i>carro forte</i>	(+73)	PT <i>arma branca</i>
8	(+79)	EN <i>noble gas</i>	(+72)	EN <i>search engine</i>
H	(−190)	PT <i>ar condicionado</i>	(−151)	PT <i>disco rígido</i>
G	(−202)	FR <i>coffre fort</i>	(−169)	EN <i>subway system</i>
F	(−202)	FR <i>bon sens</i>	(−190)	PT <i>carro forte</i>
E	(−234)	PT <i>prato feito</i>	(−238)	FR <i>disque dur</i>
D	(−292)	FR <i>baie vitrée</i>	(−256)	EN <i>half sister</i>
C	(−327)	PT <i>carta aberta</i>	(−260)	PT <i>carta aberta</i>
B	(−370)	PT <i>vinho tinto</i>	(−266)	FR <i>bonne pratique</i>
A	(−376)	PT <i>circuito integrado</i>	(−370)	EN <i>end user</i>

Table 5.7.: Outliers regarding positive/negative *geom* improvement.

5.5. Sanity checks

The number of possible DSM configurations grows exponentially with the number of internal variables in a DSM, forestalling the possibility of an exhaustive search for every possible parameter. We have evaluated above the set of variables that are most often manually tuned in the literature, but a reasonable question would be whether these results can be further improved through the modification of some other often-ignored model-specific parameters. We thus perform some sanity checks through a local search of such parameters around the highest-Spearman configuration of each DSM.

Section 5.5.1 evaluates the number of DSM iterations. Section 5.5.2 evaluates the minimum word-count threshold in the DSM. Section 5.5.3 considers a `WINDOWSIZE=2+2`. Section 5.5.4 considers higher numbers of DSM vector dimensions. Section 5.5.5 evaluates the non-determinism of DSMs through multiple random initializations. Finally, Section 5.5.6 considers whether the filtering of dataset annotations could improve its quality as well as the accuracy of predictions.

5.5.1. Number of iterations

Some of the DSMs in consideration on this chapter are iterative: they re-read and re-process the same corpus multiple times. For those DSMs, we present the results of running their best configuration, but using a higher number of iterations. This higher number of iterations is inspired by the models found in parts of the literature, where e.g. the number of *glove* iterations can be as high as 50 (Salle, Villavicencio, and Idiart 2016) or even 100 (Pennington, Socher, and Manning 2014). The intuition is that most models will lose some information (due to their probabilistic sampling), which could be regained at the cost of a higher number of iterations.

Table 5.8 presents a comparison between the baseline ρ for 15 iterations and the ρ obtained when 100 iterations are performed. For all DSMs, we see that the increase in the number of iterations does not improve the quality of the vectors, with the relatively small number of 15 iterations yielding better results. This may suggest that a small number of iterations can already sample enough distributional information, with further iterations accruing additional noise from low-frequency words. The extra number of iterations could also be responsible for overfitting of the DSM to represent particularities of the corpus, which would reduce the quality of the underlying vectors. Given the extra cost of running more iterations⁸, we refrain from building further models with as many iterations in this thesis.

5.5.2. Minimum count threshold

Minimum-count thresholds are often neglected in the literature, where a default configuration of 0, 1 or 5 being presumably used by most authors. An exception to this trend is the threshold of 100 occurrences used by Levy, Goldberg, and Dagan (2015), whose toolkit we use in *PPMI-SVD*. No explicit justification has been found for this higher word-count threshold. A reasonable hypothesis would be that higher thresholds improve the quality of the data, as it filters rare words more aggressively.

⁸ The running time grows linearly with the number of iterations.

Model (<i>FR-comp</i>)	ρ_{base}	$\rho_{\text{iter}=100}$	Difference (%)
<i>w2v-cbow</i>	.660	.640	(−2.0)
<i>w2v-sg</i>	.672	.636	(−3.7)
<i>glove</i>	.680	.677	(−0.3)
<i>lexvec</i>	.677	.671	(−0.6)
Model (<i>Reddy</i>)	ρ_{base}	$\rho_{\text{iter}=100}$	Difference (%)
<i>w2v-cbow</i>	.809	.766	(−4.3)
<i>w2v-sg</i>	.821	.777	(−4.4)
<i>glove</i>	.764	.746	(−1.8)
<i>lexvec</i>	.774	.757	(−1.7)
Model (<i>PT-comp</i>)	ρ_{base}	$\rho_{\text{iter}=100}$	Difference (%)
<i>w2v-cbow</i>	.588	.558	(−3.0)
<i>w2v-sg</i>	.586	.551	(−3.6)
<i>glove</i>	.555	.464	(−9.1)
<i>lexvec</i>	.570	.561	(−0.9)

Table 5.8.: Results using a higher number of iterations.

Table 5.9 presents the result from the highest-Spearman configurations alongside the results for an identical configuration with a higher occurrence threshold of 50. The results unanimously agree that a higher threshold does not contribute to the removal of any extra noise. In particular, for *PPMI-SVD*, it seems to discard enough useful information to considerably reduce the quality of the compositionality prediction measure. The results strongly contradict the default configuration used for *PPMI-SVD*, suggesting that a lower word-count threshold might yield better results for this task.

5.5.3. Windows of size 2+2

For many models, the best window size found was either $\text{WINDOWSIZE}=1+1$ or $\text{WINDOWSIZE}=4+4$ (see Section 5.2.1). It is possible that a higher score could be obtained by a configuration in between. While a full exhaustive search would be the ideal solution, a useful approximation of the best 2+2 configuration could be obtained by running the experiments on the highest-Spearman configurations, with the window size replaced by 2+2.

Results in Table 5.10 for a window size of 2+2 are consistently worse than the base model, indicating that the optimal configuration is likely the one that was obtained with window size of 1+1 or 4+4. This is further confirmed by the fact that most DSMs had the best configuration with window size of 1+1 or 8+8, with few cases of 4+4 as best model, which suggests that the quality of most configurations in the space of models is either monotonically increasing or decreasing with regards to these window sizes, favoring thus the configurations with more extreme WINDOWSIZE parameters.

Model (<i>FR-comp</i>)	ρ_{base}	$\rho_{\text{mincount}=50}$	Difference (%)
<i>w2v-cbow</i>	.660	.610	(−5.0)
<i>w2v-sg</i>	.672	.613	(−5.9)
<i>glove</i>	.680	.673	(−0.7)
<i>PPMI-SVD</i>	.584	.258	(−32.6)
<i>lexvec</i>	.677	.653	(−2.4)
Model (<i>Reddy</i>)	ρ_{base}	$\rho_{\text{mincount}=50}$	Difference (%)
<i>w2v-cbow</i>	.809	.778	(−3.1)
<i>w2v-sg</i>	.821	.776	(−4.5)
<i>glove</i>	.764	.672	(−9.2)
<i>PPMI-SVD</i>	.743	.515	(−22.8)
<i>lexvec</i>	.774	.738	(−3.6)
Model (<i>PT-comp</i>)	ρ_{base}	$\rho_{\text{mincount}=50}$	Difference (%)
<i>w2v-cbow</i>	.588	.580	(−0.8)
<i>w2v-sg</i>	.586	.575	(−1.1)
<i>glove</i>	.555	.540	(−1.5)
<i>PPMI-SVD</i>	.530	.418	(−11.1)
<i>lexvec</i>	.570	.566	(−0.4)

Table 5.9.: Results for a higher minimum threshold of word count.

5.5.4. Higher number of dimensions

As seen in Section 5.2.2, some DSMs obtain better results when moving from 250 to 500 dimensions, and this trend continues when moving to 750 dimensions. This behavior is notably stronger for *PPMI-thresh*, which suggests that an even higher number of dimensions could have better predictive power.

Table 5.11 presents the result of running *PPMI-thresh* for increasing values of the DIMENSION parameter. The baseline configuration (indicated as * in Table 5.11) was the highest-scoring configuration found in Section 5.2.2: *lemma_{POS}.w₁.d₇₅₀* for *PT-comp* and *FR-comp*, and *surface.w₈.d₇₅₀* for *Reddy*. As seen in Section 5.2.2, results for 250 and 500 dimensions have lower scores than the results for 750 dimensions. Results for 1000 dimensions were mixed: they are slightly worse for *FR-comp* and *Reddy*⁺⁺, and slightly better for *PT-comp*. Increasing the number of dimensions generates models that are progressively worse. These results suggests that the maximum vector quality is achieved between 750 and 1000 dimensions.

5.5.5. Random initialization

The word vectors generated by the *glove* and *w2v* models have some level of non-determinism caused by random initialization and random sampling techniques. A reasonable concern would be whether the results presented for different pa-

Model (<i>FR-comp</i>)	ρ_{base}	$\rho_{\text{win}=2+2}$	Difference (%)
<i>PPMI-SVD</i>	.584	.397	(−18.7)
<i>PPMI-thresh</i>	.702	.678	(−2.4)
<i>glove</i>	.680	.657	(−2.3)
<i>lexvec</i>	.677	.671	(−0.6)
<i>w2v-cbow</i>	.660	.644	(−1.6)
<i>w2v-sg</i>	.672	.639	(−3.3)
Model (<i>Reddy</i>)	ρ_{base}	$\rho_{\text{win}=2+2}$	Difference (%)
<i>PPMI-SVD</i>	.743	.583	(−16.0)
<i>lexvec</i>	.774	.757	(−1.7)
<i>w2v-cbow</i>	.809	.777	(−3.2)
<i>w2v-sg</i>	.821	.784	(−3.7)
Model (<i>PT-comp</i>)	ρ_{base}	$\rho_{\text{win}=2+2}$	Difference (%)
<i>PPMI-SVD</i>	.530	.446	(−8.4)
<i>PPMI-thresh</i>	.602	.561	(−4.1)
<i>lexvec</i>	.570	.564	(−0.6)

Table 5.10.: Results using a window of size 2+2.

parameter variations are close enough to the scores obtained by an average model. To assess the variability of these models, we evaluated 3 different runs of every DSM configuration (the original execution ρ_1 , used elsewhere in this thesis, along with two other executions ρ_2 and ρ_3) for *glove*, *w2v-cbow* and *w2v-sg*. We then calculate the average ρ_{avg} of these 3 executions for every model.

Table 5.12 reports the highest-Spearman configurations of ρ_{avg} for the *Reddy* and *Reddy⁺⁺* datasets. When comparing ρ_{avg} to the results of the original execution ρ_1 , we see that the variability in the different executions of the same configuration is minimal. This is further confirmed by the low sample standard deviation⁹ obtained from the scores of the 3 executions. Given the high stability of these models, results in the rest of the thesis were calculated and reported as ρ_1 for all datasets.

5.5.6. Data filtering

Along with the verification of parameters, we also evaluate whether dataset variations could yield better results. In particular, we consider the use of filtering techniques, which are used in the literature as a method of guaranteeing dataset quality. As per Roller, Schulte im Walde, and Scheible (2013), we consider two strategies of data removal: (1) removing individual outlier compositionality judgments through *z*-score filtering; and (2) removing all annotations from outlier

⁹ The low standard deviation is not a unique property of high-ranking configurations: The average of deviations for all models was .004 for *Reddy⁺⁺* and .006 for *Reddy*.

Model (<i>FR-comp</i>)	$\rho_{\text{dim}=X}$	Difference (%)
dim = 250	.671	(−3.1)
dim = 500	.695	(−0.7)
dim = 750	.702 *	(0.0)
dim = 1000	.694	(−0.8)
dim = 2000	.645	(−5.8)
dim = 5000	.636	(−6.7)
dim = 30000	.552	(−15.1)
dim = 999999	.539	(−16.3)
Model (<i>Reddy</i>)	$\rho_{\text{dim}=X}$	Difference (%)
dim = 250	.764	(−2.7)
dim = 500	.782	(−1.0)
dim = 750	.791 *	(0.0)
dim = 1000	.784	(−0.7)
dim = 2000	.760	(−3.1)
dim = 5000	.744	(−4.7)
dim = 30000	.700	(−9.1)
dim = 999999	.566	(−22.5)
Model (<i>PT-comp</i>)	$\rho_{\text{dim}=X}$	Difference (%)
dim = 250	.543	(−5.9)
dim = 500	.546	(−5.6)
dim = 750	.602*	(0.0)
dim = 1000	.609	(+0.7)
dim = 2000	.601	(−0.1)
dim = 5000	.505	(−9.7)
dim = 30000	.532	(−7.0)
dim = 999999	.500	(−10.2)

Table 5.11.: Results for higher numbers of dimensions (*PPMI-thresh*).

Dataset	DSM	configuration	ρ_1	ρ_2	ρ_3	ρ_{avg}	stddev
<i>Reddy</i>	<i>glove</i>	<i>lemma</i> _{POS} .w ₈ .d ₂₅₀	.759	.760	.753	.757	.004
	<i>w2v-cbow</i>	<i>surface</i> .w ₁ .d ₅₀₀	.796	.807	.799	.801	.006
	<i>w2v-sg</i>	<i>surface</i> .w ₁ .d ₇₅₀	.812	.788	.812	.804	.014
<i>Reddy</i> ⁺⁺	<i>glove</i>	<i>lemma</i> _{POS} .w ₈ .d ₅₀₀	.651	.646	.650	.649	.003
	<i>w2v-cbow</i>	<i>surface</i> ⁺ .w ₁ .d ₇₅₀	.730	.732	.728	.730	.002
	<i>w2v-sg</i>	<i>surface</i> ⁺ .w ₁ .d ₇₅₀	.741	.732	.721	.731	.010

Table 5.12.: Configurations with highest ρ_{avg} for non-deterministic models.

human judges. A compositionality judgment is considered an outlier if it stands at more than z standard deviations away from the mean; a human judge is deemed an outlier if its Spearman correlation to the average of the others ρ_{oth}

is lower than a given threshold R ¹⁰. These methods allow us to remove accidentally erroneous annotations, as well as annotators whose response deviated too much from the mean (in particular spammers and non-native speakers).

Table 5.13 presents the evaluation of raw and filtered datasets regarding two quality measures: the average of the standard deviations for all NCs ($\overline{\sigma_{WC}}$); and the proportion of NCs in the dataset whose standard deviation is higher than 1.5 ($P_{\sigma>1.5}$), as per Reddy, McCarthy, and Manandhar (2011). The results suggest that filtering techniques can improve the overall quality of the datasets, as seen in the reduction of the proportion of NCs with high standard deviation, as well as in the reduction of the average standard deviation itself. We additionally present the data retention rate (DRR), which is the proportion of NCs that remained in the dataset after filtering. While the DRR does indicate a reduction in the amount of data, this reduction may be considered acceptable in light of the improvement suggested by the quality measures.

Dataset	$\overline{\sigma_{WC}}$		$P_{\sigma>1.5}$		DRR
	raw	filtered	raw	filtered	
<i>FR-comp</i>	1.15	0.94	22.78%	13.89%	87.34%
<i>PT-comp</i>	1.22	1.00	14.44%	6.11%	87.81%
<i>EN-comp</i> ₉₀	1.17	0.87	18.89%	3.33%	83.61%
<i>Reddy</i>	0.99	—	5.56%	—	—

Table 5.13.: Intrinsic quality measures for the raw and filtered datasets

On a more detailed analysis, we have verified that the improvement in these quality measures is heavily tied to the use of z -score filtering, with similar results obtained when it is considered alone. The application of R -filtering by itself, on the other hand, did not show any noticeable improvement in the quality measures for reasonable amounts of DRR. This is the opposite from what was found by Roller, Schulte im Walde, and Scheible (2013) on their German dataset, where only R -filtering was found to improve results under these quality measures. We present our findings in more detail in Cordeiro, Ramisch, and Villavicencio (2016a).

We then consider whether filtering can have an impact on the performance of predicted compositionality scores. As z -score filtering was responsible for improvement in quality measures above, we consider For each of the 228 model configurations that were constructed for each language, we launched an evaluation on the filtered *EN-comp*₉₀, *FR-comp* and *PT-comp* datasets (use use z -score filtering only, as it was responsible for most of the improvement in quality measures). Overall, no improvement was observed in the results of the prediction (values of Spearman ρ) when we compare raw and filtered datasets. Looking

¹⁰ The judgment threshold we adopted was $z = 2.2$ for *EN-comp*₉₀, $z = 2.2$ for *PT-comp* and $z = 2.5$ for *FR-comp*. The human judge threshold was $R = 0.5$.

more specifically at the best configurations for each DSM (Table 5.14), we can see that most results do not significantly change when the evaluation is performed on the raw or filtered datasets. This suggests that the amount of judgments collected for each compound greatly offsets any irregularity caused by outliers, making the use of filtering techniques superfluous.

Dataset	<i>EN-comp₉₀</i>		<i>FR-comp</i>		<i>PT-comp</i>	
	raw	filtered	raw	filtered	raw	filtered
<i>PPMI-SVD</i>	.604	.601	.584	.579	.530	.526
<i>PPMI-TopK</i>	.564	.571	.550	.545	.519	.516
<i>PPMI-thresh</i>	.602	.607	.702	.700	.602	.601
<i>glove</i>	.538	.544	.680	.676	.555	.552
<i>lexvec</i>	.567	.572	.677	.676	.570	.568
<i>w2v-cbow</i>	.669	.665	.651	.651	.588	.587
<i>w2v-sg</i>	.665	.661	.653	.654	.586	.584

Table 5.14.: Extrinsic quality measures for the raw and filtered datasets

5.6. Error analysis

In the previous sections, we have studied the performance of the compositionality prediction framework in terms of the correlation between system predictions and human judgments. We now investigate the system output with regards to other variables that may have an impact on results, such as corpus frequency and conventionalization. We also compare the predicted compositionality scores with some patterns we previously found in human scores (see Section 3.2).

5.6.1. Frequency and compositionality prediction

Results from an evaluation of the hypothesis $h_{\text{idiom}} \approx \text{distr.freq}$ in Section 3.2.4 show that the frequency of NCs in large corpora is somewhat associated with the compositionality scores assigned by humans. We investigate whether this correlation also holds true to system predictions: are the most frequent NCs being predicted as more compositional?

In this experiment, we focus on a cross-language analysis with the *ALL-comp* dataset, which combines the $3 \times 180 = 540$ NCs from the three datasets presented in Chapter 3. Figure 5.11 presents the 540 NCs, ordered according to corpus frequency and grouped into 18 bins of 30 NCs each.¹¹ The height of each bin indicates the average of the scores predicted (using the *uniform* strategy) by a given system to the 30 NCs therein. There is a high variability in the level

¹¹ We use binning so as to smooth over the outliers.

of correlation between the corpus frequency of compounds and the prediction of the models. The level of correlation ranged from $\rho = .28$ for *PPMI-TopK* (not shown here) to $\rho = .68$ for *glove*, with the intermediate results of $\rho = .36$ for *PPMI-SVD*, $\rho = .46$ for *PPMI-thresh*, $\rho = .50$ for *w2v-sg*, $\rho = .51$ for *w2v-cbow* and $\rho = .54$ for *lexvec*. For every system, the correlation was significant ($p < 0.05$). This is in line with human judgments of compositionality, which also had a positive correlation with the frequency of the NCs.

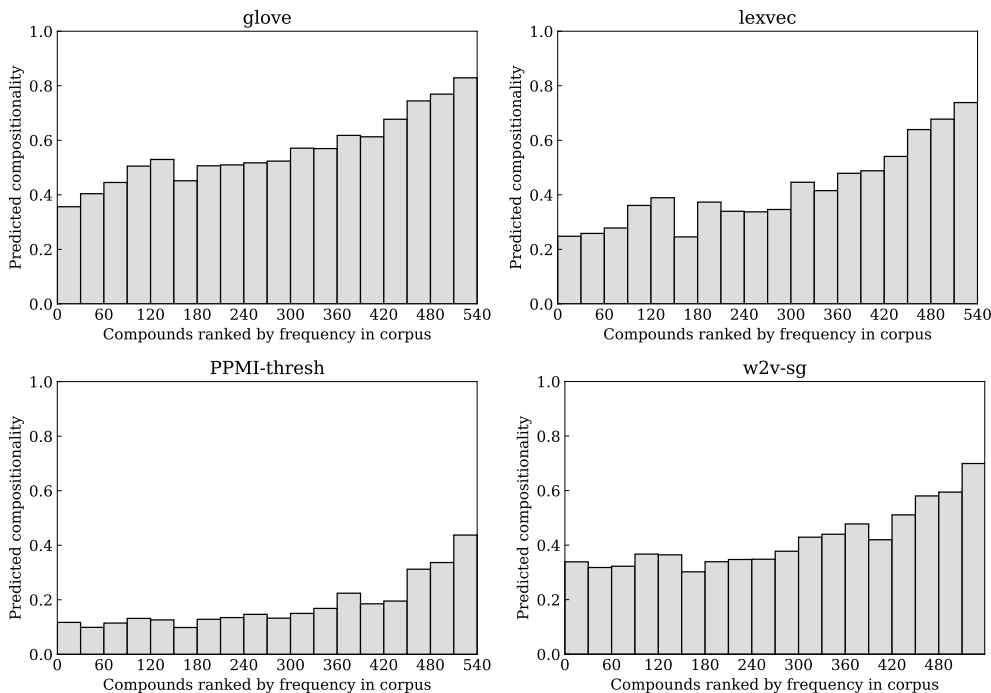


Figure 5.11.: Compositionality prediction under different frequency bins.

Another hypothesis ($h_{\text{accur}} \leftarrow \text{MWE.freq}$) we test is whether higher-frequency NCs are easier to predict. A first intuition would be that this hypothesis is true, as a higher number of occurrences is also associated with a larger amount of data, from which more representative vectors could be built. To test this hypothesis, we calculated the correlation between NC frequency and the human-system difference $|h - s|$, where h is the human score and s is the system's predicted score for a given compound. Higher values of human-system difference indicate that an NC's compositionality is harder to predict. We found a weak (though statistically significant) correlation for some of the systems: *PPMI-TopK* had $\rho = .15$, *PPMI-SVD* had $\rho = .17$, and *PPMI-thresh* had $\rho = .22$ (all with $p < 0.05$). This correlation is positive, which means that the frequency is correlated with difficulty. This implies that the compositionality of rarer NCs was mildly *easier* to predict for these systems, suggesting that the hypothesis above is false. On the other hand, *glove* had an easier time predicting frequent NC, with negative correlation of $\rho = -.19$, favoring the aforementioned hypothesis. Moreover, the correla-

tion was not statistically significant for *lexvec* and *w2v* models. These results are mixed, and either point to an overall lack of correlation between frequency and difficulty, or indicate mild DSM-specific behaviors, which should be investigated in further research.

5.6.2. Conventionalization and compositionality prediction

Section 2.2.3 has described the PMI as one well-known estimator of the level of MWE conventionalization. Many of the DSMs investigated on this thesis also rely on PMI as a way to estimate the strength of association between two words. This measure is then directly applied to target–context word pairs during the construction of the DSM, and the result becomes an internal matrix that is further processed to build the real-valued output vectors. In light of the results from the previous section, and given the reliance of most DSMs on the PMI for the construction of their word representation, one might expect similarly high correlations between compositionality and the PMI of compound elements. On the other hand, given the lack of correlation found between the conventionalization and human judgments of compositionality, good system predictions should ideally not correlate with measures of conventionalization such as the PMI. We thus evaluate whether our model really is predicting something different from conventionalization.

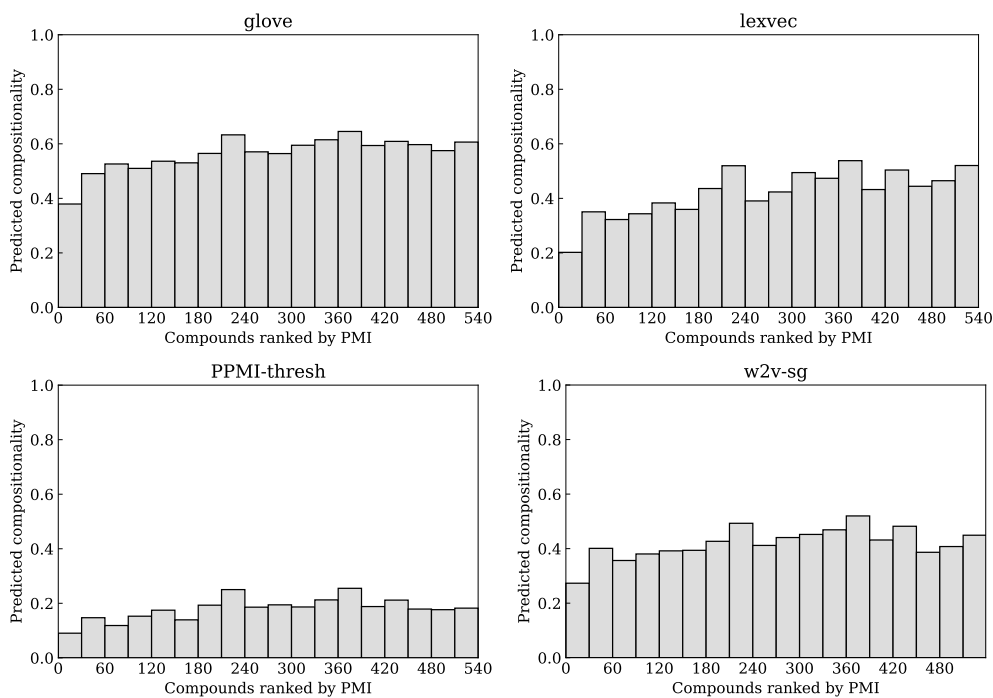


Figure 5.12.: Compositionality prediction under different PMI bins.

Figure 5.12 presents the 540 NCs of *ALL-comp*, ordered according to PMI and

grouped under 18 bins of 30 NCs each. The height of each bin indicates the average of the scores predicted (using the *uniform* strategy) by a given system to the 30 NCs therein. The effects are milder than the ones seen for the frequency (in Section 5.6.1). Statistically significant correlations are $\rho = .13$ for *PPMI-thresh*, $\rho = .17$ for *w2v-cbow* and *w2v-sg*, $\rho = .26$ for *glove* and *lexvec*, and $\rho = .28$ for *PPMI-SVD*. No correlation was found for *PPMI-TopK*. Overall, these results suggest that the vector representations generated by these models preserve some level of information regarding the strength of association between words. Given that there was no correlation between PMI and human-rated compositionality when testing hypothesis $h_{\text{idiom}} \approx \text{distr.convent}$ in Section 3.2.4, the systems that do keep this information are at a disadvantage. Particularly in the case of *w2v* models, this result is surprising, as it suggests that its high scores could be further improved by a method that did not keep as much of a correlation with the PMI in the word-embedding representation.

We also calculated the correlation between the PMI and the human–system difference, calculated as $|h - s|$, where h is the human score and s is the predicted system score for a given NC. The hypothesis ($h_{\text{accur}} \leftarrow \text{MWE.convent}$) is that the DSMs should have lower accuracy when dealing with less conventionalized NCs (and whose elements are not strongly associated through PMI), due to a lower amount of shared contexts. However, for almost all DSMs, the results obtained do not show a statistically significant correlation, suggesting that this hypothesis is not true. For *lexvec*, there was a minor negative correlation of $\rho = -.12$ ($p < 0.05$) between the PMI and the difficulty, indicating that NCs with higher PMI do have slightly more accurate internal representation than the others in this particular DSM. This differs from the results obtained when comparing the human–system difference with NC frequency (Section 5.6.1), in which *lexvec* did not show any statistically significant correlation, but most other models did. As in the case of frequency, the *w2v* models showed no correlation between difficulty of prediction and the PMI.

5.6.3. Human–system comparison

The general hypothesis $h_{\text{pred-comp}} \approx \text{comp}$ predicted a correlation between human-rated NC compositionality and model predictions, and this has been extensively verified in the highest-Spearman predictions (e.g. in Section 5.1). In this section, we present a visual validation of this hypothesis, by considering the highest-Spearman predictions of 4 DSMs, with all datasets combined.

Figure 5.13 presents 4 graphs (one per DSM), with the predicted compositionality of the NCs in *ALL-comp* for the best configuration of each language. The NCs were ranked by human compositionality scores, and grouped under 18 bins of 30 NCs each. The height of each bin indicates the average of the scores predicted (using the *uniform* strategy) by a given system to the 30 NCs therein. The four systems present a behavior that is consistent with their Spearman scores

(see Section 5.1), where system predictions grow along with the corresponding human ratings.

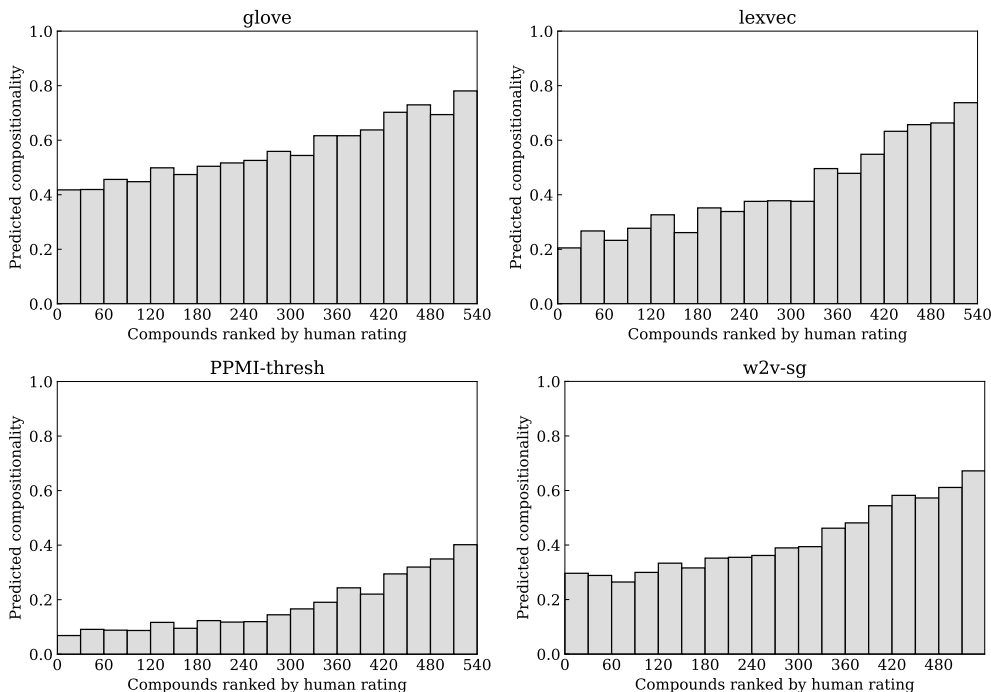


Figure 5.13.: Compositionality prediction as a function of human judgments.

For all systems but *PPMI-thresh*, the pattern of predicted compositionality grows mostly linearly with respect to the human scores (and this includes *PPMI-TopK* and *PPMI-SVD*, not shown here). The *PPMI-thresh* system ratings behave unusually, with overall lower predicted scores and a super-linear pattern of predictions, suggesting that the model is quite capable of capturing different levels of compositionality for the most compositional NCs, but fails at capturing the compositionality on the idiomatic side of the spectrum. This pattern may be explained by the fact that *PPMI-thresh* uses a sparse context representation (without any kind of dimensionality reduction other than context filtering), which means that the intersection of two vectors is often a vector with zero in many dimensions, yielding overall lower scores, especially for more idiomatic cases.

5.6.4. Range-based analyses

The Spearman score assesses the performance of a given model by providing a single numerical value. This facilitates the comparison between different models, but it hides the internal behavior of the predictions. By splitting the datasets into different ranges, we obtain a more fine-grained view of the pattern that governs

the prediction of each model.¹²

Figure 5.14 presents the highest-Spearman models (seen in Section 5.1), evaluated separately on 3 different sub-datasets of 60 NCs, split according to the standard deviation among human annotators (low, mid-range, and high values of σ_{WC})¹³. High values of standard deviation indicate disagreement among annotators, which can be regarded as an indicator that the annotation was difficult for humans. We can see that low-deviation NCs obtained considerably better system scores than the NCs for which humans disagreed among themselves. This can be taken as an evidence in favor of the hypothesis $h_{\text{accur}} \leftarrow \text{MWE.diffic}$ that higher scores are achieved for NCs that were easier for humans to annotate (i.e. that had lower standard deviation of human ratings), and suggests that part of the difficulty of this task is related to the inability of humans to determine a consensual interpretation for each NC.




Figure 5.14.: Spearman of best *uniform* models, separated by σ_{WC} ranges.

We have similarly evaluated the datasets based on three ranges of compositionality scores (low, mid-range and high values of c_{WC}). The underlying hypothesis ($h_{\text{accur}} \leftarrow \text{MWE.idiom}$) was that compositional NCs would be more precisely classified

¹² The experiments in this section involve *Reddy++*, *FR-comp* and *PT-comp*, but not *Farahmand*, as the latter dataset has binary judgments and thus cannot be easily split in ranges.

¹³ All Spearman scores for sub-datasets had $p < 0.05$.

by the model than idiomatic NCs, as the former have been more extensively considered in the literature (Mitchell and Lapata 2010; Mikolov, Sutskever, Chen, et al. 2013). Here, we consider the 540 NCs of *ALL-comp*, divided in three sub-datasets based on the level of human-rated compositionality, with 180 NCs in each sub-dataset¹⁴. Table 5.15 presents the Spearman score obtained on each sub-dataset for the highest-Spearman configuration of each DSM.¹⁵ The results suggest that distinctions on the level of compositionality are easier to perform for compositional compounds than they are for idiomatic compounds. In all cases, however, the result for sub-datasets was far lower than the score obtained for the full dataset. This might be explained by the fact that it is harder to make fine-grained distinctions of compositionality, while inter-range distinctions are more straightforward. In other words, it is easier to distinguish between an idiomatic compound (such as *ivory tower*) and a compositional one (such as *access road*) than it is to distinguish between two compositional compounds (such as *access road* and *subway system*).

Model	full dataset	low	mid	high
<i>PPMI-thresh</i>	0.66	0.29	0.24	0.37
<i>glove</i>	0.63	0.27	0.26	0.35
<i>lexvec</i>	0.64	0.18	0.20	0.37
<i>w2v-sg</i>	0.66	0.16	0.24	0.32

Table 5.15.: Spearman of best *uniform* models, separated by c_{WC} ranges.

The results above suggest that higher scores could be obtained by considering only the compounds with scores in the two extremities: lowest and highest compositionality. We evaluate this hypothesis for a given DSM by merging the predictions of its highest-Spearman configurations for *Reddy*⁺⁺, *FR-comp* and *PT-comp* (creating a single set of 540 compositionality predictions). We then consider different subsets of NC predictions in the extremities. In particular, for every window w from 1 to $270 = 540/2$, we consider the subset of w NC with lowest score prediction along with the subset of w NCs with the highest score prediction. We then calculate the Spearman ρ for this subset of $2w$ NCs, for different values of w . Figure 5.15 presents such results. As can be seen, for all 4 DSMs considered, low values of w consistently result in high Spearman scores, suggesting that the DSMs encode enough semantic information to make coarse-grained distinctions of compositionality. As we consider increasingly more cases of partially-compositional NCs (with higher values of w), we obtain increasingly lower results, until we arrive at the whole dataset of 540 NCs, where we get the lowest Spearman scores in every DSMs.

We have additionally performed both standard-deviation and compositionality-

¹⁴ Scores from compounds in different languages are mixed together in each sub-dataset.

¹⁵ All Spearman scores for datasets and sub-datasets had $p < 0.05$.




Figure 5.15.: Compositionality sliding windows, evaluating top w + bottom w compounds, for different values of w .

range analyses for other prediction strategies than *uniform*. In the case of *arith*, the Spearman score for different sub-datasets followed very closely the results of *uniform*. In the case of *maxsim*, we hypothesized that its favoring of a compositional reading of every compound would optimize results for the compositional sub-dataset when compared to *uniform*. Nevertheless, the results fluctuated around the *uniform* scores, with no clear pattern of improvement for this model. As for *geom*, we previously hypothesized that their tendency to lowering the compositionality score would optimize the quality of prediction for idiomatic compounds. The results refuted this hypothesis. Most scores were similar to *uniform* scores, with improvements seen more often in the compositional range than in the idiomatic range. However, even then the differences were small and the pattern of improvement unclear.

5.7. Summary and discussion

In this chapter, we have described the results of a large-scale evaluation of parameter choices in a DSM-based framework of compositionality prediction. Evaluations were performed on six datasets, spanning across three languages. We have built 228 DSMs for each language, and evaluated more than 8 thousand prediction model configurations, examining the impact of DSM choice and various types of parameters.

The compositionality prediction model proposed in this thesis was implemented as part of the mwetoolkit, and is freely available online.¹⁶ Given the large amount of experiments performed in this thesis, and in order to guarantee the reproducibility of results, we defined our experiments through a system of file dependencies. Every step of preprocessing (e.g. re-tokenization of compounds as a single unit, removal of stopwords) was defined in term of these dependencies, so that any modification in the code (e.g. bug fixes) would automatically invalidate experiment results. The results presented in this chapter were obtained under this system of dependencies.

Considering the experimental results in terms of DSMs, the *w2v* models performed better than *PPMI* for *Reddy*⁺⁺, both were in a tie for *Farahmand*, and *w2v* was outperformed by *PPMI-thresh* for *FR-comp* and *PT-comp*. The performance of *glove* on English datasets was underwhelming, and might be related to the lack of tuning of model-specific parameters. As previously argued by Salehi, Cook, and Baldwin (2015), *PPMI-TopK* is not an appropriate DSM for this task, as it does not model relevant co-occurrence very well.

When comparing DIMENSION across languages and datasets, larger values often bring better performance, likely due to the possibility of representing more fine-grained semantic distinctions (in agreement with the hypothesis $h_{\text{accur}} \leftarrow \text{DSM.dims}$). An upper limit of around 1000 dimensions has been verified, however, with even higher numbers of dimensions obtaining lower scores.

The most effective WINDOWSIZE depends on the model and language, but for the best models in all datasets, a window of 1+1 outperforms the others (which suggests that $h_{\text{accur}} \leftarrow \text{DSM.window}$ is false). This may be a consequence of the fact that higher window sizes are more likely to consider unrelated words as part of a target’s context.

Regarding the WORDFORM, the *lemma* (i.e. stopword removal + lemmatization) seems to be the overall best type of preprocessing across languages (as predicted by $h_{\text{accur}} \leftarrow \text{corpus.wordform}$). The use of POS tags does not seem to improve on the results, which could indicate that the higher precision of grammatical category does not compensate for the added sparsity. In the case of English, the effects of both stopword removal and lemmatization are questionable, with plain surface-level word-forms producing slightly better models in some cases.

Corpus size seems to play a fundamental role in the quality of the constructed distributional models, as corpora with less than a billion tokens result in considerably weaker predictions ($h_{\text{accur}} \leftarrow \text{corpus.size}$). However, the improvement in prediction quality seems to be capped at around a threshold of one billion tokens: larger corpora do not result in better predictions of compositionality for nominal compounds. This threshold may be related to the minimum frequency necessary for rarer NCs so as to permit the calculation of cosine similarity with its components.

¹⁶ <http://mwetoolkit.sf.net>

The technique of parallel predictions was shown to perform equivalently to whole-corpus predictions ($h_{\text{accur}} \leftarrow \text{corpus.parallel}$). While the use of this technique does not improve on the results obtained through whole-corpus models, it does permit a more flexible utilization of computational resources (e.g. clusters) in the construction of the underlying semantic representations.

Regarding the different compositionality prediction strategies, the *uniform* strategy produces predictions that are consistently among the best ones. The *maxsim* strategy does improve the prediction of compositional NCs, but only for outlier cases, contributing to random variation in most cases ($h_{\text{strat.maxsim}}$). While this does not improve on the results from *uniform*, it does consistently produce similarly good results. The *head* and *mod* strategies perform surprisingly well for all top models of every dataset, in spite of their reliance on incomplete information ($h_{\text{strat.partial-info}}$). The performance of *arith* is quite similar to *uniform*, reflecting the fact that both rely on an additive model of compositionality ($h_{\text{strat.arith}} \approx \text{strat.uniform}$). The *geom* strategy did optimize the scores of idiomatic NCs, but at the expense of a pessimization of scores for some compositional cases ($h_{\text{strat.geom}}$). A combination of the *geom* and *maxsim* strategies is left for future work.

Concerning the sanity checks, we found no advantage in the use of a higher number of iterations for the construction of DSMs. The minimum word-count has similarly been found to be a small value, with higher thresholds removing too much information. An evaluation of the random initialization used in some DSMs found no difference in the final results across multiple executions. Regarding the dataset scores, filtering techniques were also considered, but the results were comparable to the ones obtained on the unfiltered datasets.

This chapter has also performed an error analysis of the predicted compositionality scores. As in the case of human-rated scores, frequency was found to be positively correlated with compositionality ($h_{\text{idiom}} \approx \text{distr.freq}$). This result disputes the hypothesis that idiomatic expressions are more frequent. In the case of PMI, while it was not correlated to human-rated scores, it did show a mild correlation with some system scores, suggesting that these systems could be improved by reducing their reliance on that measure ($h_{\text{idiom}} \approx \text{distr.convent}$). Intra-NC standard deviation on human ratings has also been shown to be related to system scores: systems have difficulty on NCs that humans also find difficult ($h_{\text{accur}} \leftarrow \text{MWE.diffic}$). Moreover, system predictions were found to have higher quality in the case of compositional expressions ($h_{\text{accur}} \leftarrow \text{MWE.idiom}$). Further work would be required to improve score predictions of idiomatic NCs.

An overall recommendation for future work would be the use of large dimensions and small window sizes. Moreover, investing in preprocessing provides a good balance of a small vocabulary (of lemmas) and good accuracy. The underlying corpus size should contain at least 1 billion tokens. As for the underlying model, the simple *uniform* prediction strategy can achieve the highest-quality predictions.

Regarding the choice of DSM, the average Spearman’s ρ for *Reddy* over all tested parameter configurations was 0.71 for both *w2v* models and 0.67 for *PPMI-thresh*, suggesting that both types of models can obtain good results. While *PPMI-thresh* is a simple, fast and inexpensive model to build, *w2v* has a free and push-button implementation, and requires less hyper-parameter tuning, as it seems more robust to parameter variation.

More generally, the best results obtained are comparable and even outperform the state of the art. Table 5.16 compares the highest results in the literature for the *Reddy* dataset against the highest-Spearman and highest-Pearson configuration obtained for each DSM.¹⁷ Reddy, McCarthy, and Manandhar (2011) use a compositionality prediction model with a global set of contexts that resembles *PPMI-TopK*, and the results are correspondingly similar to the ones obtained for this DSM. Salehi, Cook, and Baldwin (2014) also use global contexts, but augment it with information obtained from translations, which improves the results (they are somewhat comparable to our highest-Pearson *PPMI-SVD* configuration). Salehi, Cook, and Baldwin (2015) use a configuration that is similar to our highest-Pearson *w2v-cbow*. We obtain slightly better results due to our exploration of the space of DSM and corpus configurations.¹⁸

Model & Parameters	Spearman ρ	Pearson r
Reddy, McCarthy, and Manandhar (2011)	.714	—
Salehi, Cook, and Baldwin (2014)	—	.744
Salehi, Cook, and Baldwin (2015)	—	.796
Best <i>w2v-sg</i> [Spearman: <i>surface.w1.d750</i>]	.812 (.812)	.814 (.814)
Best <i>PPMI-thresh</i> [Spearman: <i>surface.w8.d750</i>]	.791 (.803)	.762 (.768)
Best <i>w2v-cbow</i> [Spearman: <i>surface+.w1.d500</i>]	.796 (.796)	.803 (.798)
Best <i>lexvec</i> [Spearman: <i>surface+.w4.d500</i>]	.774 (.773)	.787 (.787)
Best <i>glove</i> [Spearman: <i>lemmaPos.w8.d250</i>]	.754 (.759)	.783 (.787)
Best <i>PPMI-SVD</i> [Spearman: <i>surface+.w1.d500</i>]	.743 (.743)	.738 (.726)
Best <i>PPMI-TopK</i> [Spearman: <i>lemmaPos.w8.d1000</i>]	.706 (.716)	.732 (.717)

Table 5.16.: Comparison of our best models with state-of-the-art results for *Reddy*. Results in parentheses for fallback evaluation.

Our results are also comparable to the state of the art regarding the *Farahmand* dataset, particularly when the fallback evaluation is adopted, as shown in Table 5.17. The predictive model of Yazdani, Farahmand, and Henderson (2015) generalizes the linear combination of word representations (such as the one used on the *uniform* strategy) so as to allow for other polynomial projections, with quadratic projections on *w2v-cbow* obtaining the highest BF₁ score

¹⁷ Due to space constraints, only the highest-Spearman configuration is shown.

¹⁸ Note that the main goal was not to beat the state of the art, but to explore the space of configurations.

of .487. We show that a DSM and corpus parameter tuning can beat the use of these more complex functions, as our best configuration for *w2v-cbow* obtains a BF_1 of .512. Future work should investigate the joint use of quadratic projections and the recommended DSM configurations from this thesis.

Model & Parameters	BF_1
Yazdani, Farahmand, and Henderson (2015)	.487
Best <i>w2v-cbow</i> [<i>lemma.w₁.d₇₅₀</i>]	.512 (.471)
Best <i>w2v-sg</i> [<i>lemma.w₄.d₅₀₀</i>]	.507 (.468)
Best <i>lexvec</i> [<i>surface.w₁.d₇₅₀</i>]	.449 (.431)
Best <i>PPMI-SVD</i> [<i>lemma.w₄.d₇₅₀</i>]	.487 (.424)
Best <i>PPMI-thresh</i> [<i>lemma.w₄.d₇₅₀</i>]	.472 (.404)
Best <i>PPMI-TopK</i> [<i>lemma.w₈.d₁₀₀₀</i>]	.435 (.376)
Best <i>glove</i> [<i>lemma_{POS}.w₈.d₇₅₀</i>]	.400 (.358)

Table 5.17.: Comparison of our best models with state-of-the-art BF_1 for *Farahmand*. Results in parentheses for fallback evaluation.

6. Extrinsic evaluation of compositionality prediction

The accurate identification of MWEs in running text is a major challenge in the general pipeline of NLP applications. The set of all possible categories of MWEs in a language can be quite diverse (Schneider, Onuffer, Kazour, et al. 2014a; Constant, Eryigit, Monti, et al. 2017), and the often-employed method of identifying such expressions from a predetermined lexicon may not yield satisfactory results for productive MWE patterns (such as nominal compounds). MWE identification has notably been one of the goals of the SemEval 2016 task 10: DiMSUM (Detecting Minimal Semantic Units and their Meanings) (Schneider, Hovy, Johannsen, et al. 2016). In this shared task, participants were expected to present a system that was able to detect and group MWEs, and to assign supersense tags to each semantic unit (MWE or single word).

In this chapter, we consider an extrinsic evaluation of predicted compositionality scores, which are adopted as features in a system of MWE identification. The hypothesis we want to evaluate is $h_{\text{pred-comp} \rightarrow \text{ident-accur}}$, which predicts that the task of MWE token identification should benefit from the use of compositionality scores. We focus on the identification of noun-based compounds (i.e. nominal compounds, including proper names and nominal compounds with prepositions, such as *chamber of commerce*). For the identification of other categories of MWEs (as well as our work on supersense tagging), we refer to the paper that describes our submission for the DiMSUM shared task (Cordeiro, Ramisch, and Villavicencio 2016c), as well as the paper on CRF-based detection of MWEs (Scholivet, Ramisch, and Cordeiro 2017). Section 6.1 presents two methods of MWE identification. Section 6.2 describes the experimental setup for the extrinsic evaluation. Section 6.3 then presents the results obtained with and without compositionality scores. Finally, Section 6.4 concludes with the summary of the main findings from this chapter.

6.1. Proposed models of MWE identification

In the interest of validating the compositionality prediction model proposed in this thesis, we consider two methods of MWE token identification, both of which can be applied with or without the feature of compositionality scores. The task of MWE identification consists in taking a tokenized corpus as input and generating an extra layer in which every occurrence of an MWE is explicitly indicated.¹

¹ A full review of MWE identification methods is out of the scope of this work. We refer to the *MWE Identification* section of Constant, Eryigit, Monti, et al. (2017) for a thorough survey of

We consider two techniques of MWE identification: a rule-based method and a probabilistic method.

6.1.1. Rule-based identification model

Rule-based methods identify MWE occurrences by projecting type-level representations from a lexicon onto a layer of MWE occurrences in a corpus (see Section 2.3.3 on MWE token identification). We propose a baseline model of rule-based MWE identification which identifies words in the corpus that correspond to MWE entries in the lexicon. This identification is based on lemmas and POS tags, and may be done on a preexisting lexicon or on a list of MWE candidates extracted through techniques of MWE type discovery (described in Section 2.3.2).

We perform MWE token identification using an augmented version of the `mwe-toolkit`, including support for both type-level discovery and token-level identification of contiguous and non-contiguous MWEs based on some degree of customization (Cordeiro, Ramisch, and Villavicencio 2015). MWE type-level candidates are extracted from a training corpus through syntactic patterns, without losing track of their token-level occurrences, to guarantee that all the MWE occurrences learned from the training data can be projected onto the test corpus. These candidates can then be filtered based on a variety of conditions (in particular, whether their occurrences are always annotated in the training corpus). The resulting set of candidates can then be automatically projected onto a layer of corpus MWE occurrences. We will use this as a baseline model, and as such, the identification will be context-independent (identifying every possible occurrence as an MWE regardless of any contextual clues).

These are the main functionalities that we have developed and integrated into the `mwe-toolkit` for experiments on MWE identification:

1. Different match distances:
 - Longest: Matches the longest possible candidate. Useful e.g. for nominal compounds, where we want to match the whole compound.
 - Shortest: Matches the shortest possible candidate. Useful e.g. for phrasal verbs, where we want to find only the closest particle.
 - All: Matches all possible candidates. Useful as a fallback when shortest and longest are too strict (post-processing is then required).
2. Different match modes:
 - Non-overlapping: Matches at most one MWE per word in the corpus.

other methods of MWE identification.

- **Overlapping:** Allows words to be part of more than one MWE. This can be used to find MWEs that occur inside the gap of another MWE, or MWE occurrences that share a token.
3. **Source-based identification:** When information is retrieved in MWE type discovery, we keep a detailed description of the source corpus and sentence. The identification step can then be quickly performed by projecting the MWEs back on the source corpus.

As an example, consider the following two MWE patterns described by regular expressions over POS tags:

- `NounCompound` \rightarrow `Noun Noun+`
- `PhrasalVerb` \rightarrow `Verb (Ignored*) Particle`

Figure 6.1 presents the results of applying different matching combinations to these patterns. Consider an input such as the one in Figure 6.1(a). By applying a non-overlapping contiguous approach to the noun compound identification and a gappy approach to the verb-particle construction, we may automatically identify two MWE candidates in the sentence. If we use the *longest* match distance for both patterns, we capture the whole nominal compound, but we go too far for the verb-particle construction (Figure 6.1(b)). The opposite happens if we use *shortest* match distance for both patterns, which works well for the verb-particle construction but does not capture the whole nominal compound (Figure 6.1(c)). By using different configurations for each type of MWE, we are able to identify the correct occurrences in the text (Figure 6.1(d)).

(a)	You	threw	those	lab	rat	tissue	samples	out	without	thinking	?
(b)	You	<u>threw</u>	those	lab	rat	tissue	samples	out	<u>without</u>	thinking	?
(c)	You	<u>threw</u>	those	lab	rat	tissue	samples	<u>out</u>	without	thinking	?
(d)	You	<u>threw</u>	those	lab	rat	tissue	samples	<u>out</u>	without	thinking	?

Figure 6.1.: MWE-annotated output with different match distances.

6.1.2. Probabilistic identification model

In addition to the rule-based approach of MWE identification, we also consider a probabilistic model, in the form of linear-chain conditional random fields (CRFs) (Lafferty, McCallum, and Pereira 2001). Under this approach, we construct a classifier that tags each input token based on whether it is independent or a part of an MWE.² The CRF is trained based on a set of observations $T = T_1 \dots T_n$, in

² The CRF tagger was trained with CRFSuite (Okazaki 2007).

which each observed input token T_i is paired up with a tag y_i . When performing predictions, the probability of a given output tag y_i for an input token T_i depends on the tag of its neighbor token (y_{i-1}), and on a set of features of the input $\phi(T)$. The feature values can come from any position of the input sequence, including the current token T_i .

We represent MWE identification as a tagging problem through the use of the Begin-Inside-Outside (BIO) encoding (Ramshaw and Marcus 1995). In a BIO representation, each token T_i in the training corpus is annotated with a corresponding tag B (beginning of the MWE), I (inside MWE) or O (independent token, outside any MWE). In this scheme, MWEs must all be contiguous, and overlaps cannot be represented.³

6.2. Experimental setup

We instantiate multiple variants of the rule-based and probabilistic methods of MWE identification. We present below the configuration that we use for each method, as well as the annotated corpora on which they are evaluated. For both methods, we consider the task identification with and without a *compositionality* feature, derived from a lexicon of predicted compositionality scores, which we also describe below.

6.2.1. Reference corpora

We perform the evaluation of the MWE identification models on two MWE-annotated corpora.

- For the English language: training, development and test data are the ones provided for the DiMSUM shared task (Schneider, Hovy, Johannsen, et al. 2016). The sentences originally come from multiple English corpora: a corpus of online reviews (STREUSLE), two corpora containing Twitter data (Ritter and Lowlands), and a corpus built from TED Talks transcripts⁴. The resulting training corpus contains 4 800 sentences, and the test corpus contains 1 000 sentences. Every MWE annotation was reviewed by at least two annotators. The authors do not report the annotator agreement (Schneider, Onuffer, Kazour, et al. 2014a; Schneider, Hovy, Johannsen, et al. 2016).
- For the French language: training, development and test data come from an adaptation of the French Treebank (FTB) (Abeillé, Clément, and Toussenet 2003) from the SPMRL shared task on Statistical Parsing of Morphologically-Rich Languages (Seddah, Tsarfaty, Kübler, et al. 2013). The corpus consists

³ Note however that more complex schemes could account for some level of discontinuity and overlap, such as the two-layer B \tilde{B} I \tilde{I} O \tilde{O} representation of Schneider (2014).

⁴ The train/dev corpora did not contain sentences from TED. Only the blind test data did.

of a collection of newspaper entries (*Le Monde*) from multiple domains, with a total of around 1 million tokens. It contains manually-validated lemmas, POS-tag annotations, syntactic information (ignored in this work) and a layer of MWE occurrence annotations.

In both corpora, we keep only MWEs representing nominal compounds. This is done through a pattern-based filtering on the MWE layer, using the *mwetoolkit*. The goal of this step is to filter out all MWEs that do not contain a noun (e.g. *by and large*), as well as MWEs that contain verbs (e.g. *give birth*). The resulting English test corpus has 254 MWEs, and the resulting French test corpus has 849 MWEs.

6.2.2. Compositionality lexicons

In both rule-based and probabilistic methods, we consider the use of a *compositionality* feature, which we derive from a lexicon of predicted compositionality scores. The lexicon itself was constructed through a type-based extraction of MWE candidates from the reference corpora. The extraction used a language-specific pattern. For English, we allow adjective+noun pairs (e.g. *red wine*) as well as noun+preposition+noun (e.g. *cup of tea*), and combinations thereof (e.g. *president of the United States*). For French, we consider noun+adjective pairs (e.g. *vin blanc* (lit. *wine white*)), adjective+noun pairs (e.g. *longue durée*, ‘long-term’ (lit. *long duration*)), as well as noun+preposition+noun expressions (e.g. *mise à jour* ‘update’ (lit. *put to day*)), including combinations of these (e.g. *Journal officiel de la République Française* ‘Official gazette of the French Republic’ (lit. *Newspaper official of the Republic French*)).

For each language, we projected the extracted MWE candidates onto WaC corpora⁵, and then constructed two DSMs instances (*w2v-sg* and *PPMI-thresh*), with the same setup as in Section 4.4, using lemma.w₁.d₇₅₀. For each DSM, we calculated the predicted compositionality score for each MWE candidate under the *uniform* strategy. This resulted in a total of four lexicons of compositionality (varying between two DSMs and two languages).

6.2.3. Rule-based identification

For each language, our baseline rule-based MWE identification algorithm considers 7 different rule configurations. Two of these rules are directly based on data from the training corpus, two are based on an approach of MWE identification based on POS-tag patterns, and two are based on the previously described compositionality lexicons (described in Section 6.2.2).

For the rules based on training data, annotated MWEs are extracted from the training corpus and then filtered. We keep MWE candidates whose proportion

⁵ We use the same corpora as in Section 4.3.

of annotated instances with respect to all occurrences in the training corpus is above a threshold τ , discarding the rest. For the selection of thresholds, we refer to Cordeiro, Ramisch, and Villavicencio (2016c), where we considered thresholds $\tau \in \{0\%, 10\%, 20\%, \dots, 100\%\}$, obtaining the best results for $\tau = 40\%$ (contiguous MWEs) and $\tau = 70\%$ (gappy MWEs).

The last step of rule-based identification consists in projecting the filtered list of MWE candidates on the test data, that is, we segment as MWEs the test token sequences that are contained in the lexicon extracted from the training data. These configurations are:

- $\text{TRAIN}_{\text{CONTIG}}$: Contiguous MWEs annotated in the training corpus at least once are extracted and filtered with a threshold of $\tau = 40\%$. That is, we create a lexicon containing all contiguous lemma+POS sequences for which at least 40% of the occurrences in the training corpus were annotated (e.g. we keep the expression *last minute*, as it was annotated in $\frac{5}{6} = 83\%$ of its occurrences in the training data). The resulting lexicon is projected on the test corpus using this rule: an MWE is deemed to occur if its component words appear contiguously in a sentence.
- $\text{TRAIN}_{\text{GAPPY}}$: Non-contiguous MWEs are extracted from the training corpus and filtered with a threshold of $\tau = 70\%$. The resulting MWEs are projected on the test corpus using the following rule: an MWE is deemed to occur if its component words appear sequentially with at most a total of 3 gap words in between them. This method is not used for French, as only contiguous MWEs were annotated in the corpus.

We also identify MWEs in the test corpus based on POS-tag patterns:

- $\text{PATTERN}_{\text{NOUN}}$: We collect candidate nominal compounds from the test corpus that never appear in the training corpus, and project them back on the test corpus. For English, we focus on contiguous noun+noun sequences (e.g. *car wash*), as they are the most prevalent in the DiMSUM corpus. For French, we consider contiguous noun+adjective pairs.⁶ As the French corpus does not distinguish common nouns from proper nouns, both are included as part of this method.
- $\text{PATTERN}_{\text{PROPN}}$: The English corpus distinguishes common nouns and proper nouns through their POS tag. In this method, we annotate sequences of two or more tokens POS tagged as proper nouns (PROPN), in an effort to identify named entities such as *New York City*. We do not consider any thresholds, as named entities are sparse and most occurrences from training do not appear in test.

⁶ Syntactic structures involving combinations of nouns, adjectives and prepositions are rarely annotated in this corpus.

For each language, we also consider two methods of MWE identification based on compositionality lexicons (Section 6.2.2). We annotate as MWE every contiguous occurrence of an entry in the compositionality lexicon, as long as its compositionality score is under a given threshold ($CS_\beta \leq \text{threshold}$). We consider thresholds between 0 (most restrictive, eliminates almost all MWEs from the lexicon) and 1 (most permissive, keeps almost all MWEs in the lexicon).

- COMP_{w2v} : Uses the compositionality lexicon built from *w2v-sg*.
- COMP_{ppmi} : Uses the compositionality lexicon built from *PPMI-thresh*.

6.2.4. Probabilistic identification

We consider the following sets of features $\phi(T)$:

- CTX: This is a set of contextual features which corresponds to the BEST_2 set from Scholivet, Ramisch, and Cordeiro (2017), without the association measures. The feature set contains 21 single-token, 2-gram and 3-gram features (involving surface-form, lemma and POS tag of tokens). It also includes features indicating: whether the current token has a hyphen, whether it has a digit, and whether it is in upper-case form. We refer to the paper for an in-depth feature analysis, as well as a broader evaluation of this model for all categories of MWEs in the French corpus.
- AM: This set of features contains four association measures: PMI, MLE, Student’s t and log-likelihood (see Section 2.2.3). These are the association measures that were found to be the most impactful in Scholivet, Ramisch, and Cordeiro (2017), when evaluating corpora with multiple categories of MWEs.
- COMP: This set of features is based on the MWE scores from the compositionality lexicons. We designate the set of features derived from *w2v-sg* as COMP_{w2v} , and the one derived from *PPMI-thresh* as COMP_{ppmi} .

Different methods of CRF modeling may or may not be able to accurately represent continuous features. In this work, we circumvent possible limitations by quantizing every numerical score (obtained in AM and COMP) using a uniform distribution; i.e. we assign an equal number of MWEs to 5 different bins based on their numerical scores. We leave the evaluation of continuous CRF models for future work (Huang, Xu, and Yu 2015).

6.3. Results

For each reference corpus, we evaluate the MWE identification models on the *development* part under a variety of setups, as described in Section 6.2. We present the results below.

6.3.1. Rule-based identification: baseline

We start with an analysis of the baseline results obtained by the rule-based MWE identifier for the English corpus. Table 6.1 presents the individual score obtained by each rule. At a first glance, the most promising rules seem to be $\text{TRAIN}_{\text{CONTIG}}$ and $\text{PATTERN}_{\text{PROP}}_{\text{N}}$, both of which obtain a high level of precision. The rule $\text{TRAIN}_{\text{GAPPY}}$ also obtains a high precision, but it does not capture many occurrences of MWEs, obtaining low recall.

Rules	Precision	Recall	F_1
$\text{TRAIN}_{\text{CONTIG}}$.843	.232	.364
$\text{TRAIN}_{\text{GAPPY}}$.750	.012	.023
$\text{PATTERN}_{\text{PROP}}_{\text{N}}$.750	.272	.399
$\text{PATTERN}_{\text{NOUN}}$.315	.181	.230

Table 6.1.: Baseline results for rule-based MWE identifier (English dataset).

We then consider the accuracy of MWE identification when multiple rules are combined. In particular, we fix the highest-ranking rule $\text{TRAIN}_{\text{CONTIG}}$, and we consider combinations involving the other rules. Table 6.2 presents the new results. The addition of the rule $\text{TRAIN}_{\text{GAPPY}}$ does manage to slightly improve the recall of $\text{TRAIN}_{\text{CONTIG}}$, but at the expense of a considerable decrease in the precision. The addition of the rule $\text{PATTERN}_{\text{PROP}}_{\text{N}}$ improves both precision and recall, reflecting the fact that named entities are sparse, and most occurrences were not seen in training data. Similarly, the addition of the rule $\text{PATTERN}_{\text{NOUN}}$ does improve recall, suggesting that many occurrences of noun+noun compounds were not seen in the training data. However, many of these predictions are spurious (i.e. they refer to productive combinations of nouns, such as *dinner plate*, which were not annotated), and thus the precision of these 2 combined rules is sub-par. The same behavior can be seen in the last line, where we consider all rules but $\text{TRAIN}_{\text{GAPPY}}$. The combination achieves a considerably higher recall, but at the expense of a reduction in the precision.

We similarly consider the baseline results obtained by the rule-based MWE identifier for the French corpus. Table 6.3 presents the individual score obtained by each rule. Differently from the English corpus, many of the MWEs in the French development set had a counterpart in the training set, which contributed to a $\text{TRAIN}_{\text{CONTIG}}$ recall of more than 60% of the occurrences. The pattern-based

Rules	Precision	Recall	F ₁
TRAIN _{CONTIG} + TRAIN _{GAPPY}	.831	.232	.363
TRAIN _{CONTIG} + PATTERN _{PROPN}	.783	.484	.599
TRAIN _{CONTIG} + PATTERN _{NOUN}	.491	.413	.449
TRAIN _{CONTIG} + PATTERN _{PROPN} + PATTERN _{NOUN}	.561	.665	.609

Table 6.2.: Combined baselines for rule-based MWE identifier (English dataset).

rule PATTERN_{NOUN} does find a modest amount of new MWEs, but at the cost of a very low precision. Therefore, when both rules are combined, the result is quite a bit lower than the one obtained for TRAIN_{CONTIG} alone.

Rules	Precision	Recall	F ₁
TRAIN _{CONTIG}	.862	.684	.763
PATTERN _{NOUN}	.081	.107	.092
TRAIN _{CONTIG} + PATTERN _{NOUN}	.381	.792	.515

Table 6.3.: Baseline results for rule-based MWE identifier (French dataset).

6.3.2. Rule-based identification: compositionality scores

The rule-based method can also be applied with an external lexicon of MWEs. We consider the two lexicons described in Section 6.2.2: COMP_{W2V} and COMP_{PPMI}. These lexicons associate MWE candidates with an automatically-calculated compositionality score. We consider multiple variants of the MWE identification model by applying different thresholds on what score constitutes an idiomatic MWE. Lower thresholds should improve precision (as they only allow annotation of highly non-compositional cases), while reducing the recall due to their restrictiveness.

Table 6.4 presents the results obtained for the rule-based system with different thresholds of compositionality scores in COMP_{W2V}. As expected, lower thresholds are associated with a lower recall in both languages. The precision, on the other hand, presents an unexpected behavior: the lowest precision (indicated through [†] on the table) is not associated with the more permissive threshold of $CS_{\beta} \leq 1.000$. In fact, in the case of the French corpus, the precision falls monotonically as we consider stricter thresholds. In the case of the English corpus, the lowest precision is associated with a middle-range threshold of $CS_{\beta} \leq 0.200$, but note that other more restrictive thresholds (indicated through [?] on the table) are much less reliable, as the number of MWEs predicted by the system is very low ($CS_{\beta} \leq 0.100$ has 37 predictions, $CS_{\beta} \leq 0.050$ has 10 predictions, and $CS_{\beta} \leq 0.000$ has 4 predictions). In the case of French, the lowest threshold of $CS_{\beta} \leq 0.000$ produces 85 predictions.

Threshold	English corpus			French corpus		
	Precision	Recall	F ₁	Precision	Recall	F ₁
$CS_{\beta} \leq 1.000$.152	.185	.167	.245	.482	.325
$CS_{\beta} \leq 0.500$.149	.181	.164	.245	.482	.325
$CS_{\beta} \leq 0.200$.098 [†]	.055	.070	.176	.239	.203
$CS_{\beta} \leq 0.100$.216 [?]	.031	.055	.132	.111	.120
$CS_{\beta} \leq 0.050$.200 [?]	.008	.015	.128	.029	.048
$CS_{\beta} \leq 0.000$.250 [?]	.004	.008	.012 [†]	.001	.002

Table 6.4.: Results for rule-based MWE identifier using comp_{w2v} .

We then consider whether a similar behavior can be observed when using compositionality scores from $\text{COMP}_{\text{PPMI}}$. Table 6.5 presents the results obtained for the rule-based system under these scores. As seen in Section 5.6.3, *PPMI-thresh* scores tend to be lower than in other DSMs, which explains why the scores obtained for high thresholds are almost identical. As in the case of COMP_{w2v} above, lower thresholds are associated with a lower recall in both languages. Moreover, the lowest values of precision are once again associated with lower thresholds, suggesting that this is a consistent property of both datasets. These precision scores are more reliable than the ones from COMP_{w2v} (for English, $CS_{\beta} \leq 0.000$ has 77 predictions, $CS_{\beta} \leq 0.005$ has 177 predictions and $CS_{\beta} \leq 0.010$ has 261 predictions; while for French, $CS_{\beta} \leq 0.000$ has 513 predictions, $CS_{\beta} \leq 0.005$ has 1648 predictions and $CS_{\beta} \leq 0.010$ has 2155 predictions).⁷

Threshold	English corpus			French corpus		
	Precision	Recall	F ₁	Precision	Recall	F ₁
$CS_{\beta} \leq 1.000$.118	.240	.158	.136	.522	.215
$CS_{\beta} \leq 0.500$.118	.240	.158	.136	.522	.215
$CS_{\beta} \leq 0.200$.118	.240	.158	.135	.519	.215
$CS_{\beta} \leq 0.100$.118	.236	.158	.134	.514	.213
$CS_{\beta} \leq 0.050$.100	.185	.130	.124	.461	.196
$CS_{\beta} \leq 0.020$.094	.138	.112	.090	.284	.137
$CS_{\beta} \leq 0.010$.077 [†]	.079	.078	.060	.153	.087
$CS_{\beta} \leq 0.005$.090	.063	.074	.053 [†]	.104	.070
$CS_{\beta} \leq 0.000$.104	.031	.048	.066	.040	.050

Table 6.5.: Results for rule-based MWE identifier using $\text{comp}_{\text{ppmi}}$.

A possible explanation to the behavior observed above would be that these datasets also annotate collocations along with idiomatic MWEs. We investigate this hypothesis through a manual annotation of all 309 MWE candidates identified by COMP_{w2v} . We classify each MWE occurrence in one of these 4 categories:

⁷ Note that, due to the nature of PPMI vs PMI, $CS_{\beta} \leq 0$ is equivalent to $CS_{\beta} = 0$ for $\text{COMP}_{\text{PPMI}}$.

- Productive expression: when the words are combined in a fully productive manner (e.g. *nice car*). In a productive expression, any of the elements can be replaced by a similar word without any loss of meaning or increased markedness (e.g. *nice car* → *nice airplane*, *cool car*).
- Collocation: when the choice of words is conventionalized, but is still compositional (e.g. *test results*). In these cases, changes in word order and replacement by similar words or synonyms is possible, but has a distinctive markedness (e.g. the expression *results from the tests* still refers to a similar concept, but is not the preferred way of referring to *test results*).
- Idiomatic expression: when the whole expression is idiomatic, with at least one of the words not contributing to a literal sense. This includes crystallized metaphors (e.g. *extra mile*) and proper names (e.g. *New Jersey*).
- Other: for all other cases. This includes errors of POS tag (e.g. *wind blows* classified as NOUN+NOUN), adjacent sequences of words that do not form a phrase (e.g. *home from work*), and cases in which the intended meaning was considered hard to judge even in the context of the original sentence (e.g. an occurrence of *good sport*).

For the first three categories, we consider the fraction of their occurrences for each threshold of predicted compositionality score (Table 6.6). In particular, we consider the lines with $CS_\beta \leq 0.5$ and $CS_\beta \leq 0.2$. In the case of productive MWEs, for these two thresholds, we can see that they present a similar rate of occurrence among the different levels of compositionality score (around 55%). The rate of occurrence of idiomatic expressions has a slight variation, with a higher rate for $CS_\beta \leq 0.2$, as expected. Note however the case of collocations: they present a difference of more than 7% between the two thresholds.⁸ This suggests that the difference between annotations is related to the fact that the dataset includes many collocations, which is precisely what we filter out when we use the compositionality scores.

Threshold	Productive	Collocation	Idiomatic
$CS_\beta \leq 1.000$	170/309 = 55.0%	61/309 = 19.7%	28/309 = 9.1%
$CS_\beta \leq 0.500$	170/308 = 55.2%	60/308 = 19.5%	28/308 = 9.1%
$CS_\beta \leq 0.200$	78/143 = 54.5%	17/143 = 11.9%	17/143 = 11.9%
$CS_\beta \leq 0.100$	17/37 = 45.9%	6/37 = 16.2%	8/37 = 21.6%
$CS_\beta \leq 0.050$	1/10 = 10.0%	3/10 = 30.0%	4/10 = 40.0%
$CS_\beta \leq 0.000$	0/4 = 0.0%	1/4 = 25.0%	3/4 = 75.0%

Table 6.6.: Classification of MWE candidates identified by comp_{w2v} .

⁸ Note that the scores for $CS_\beta \leq 0.1$ and for lower thresholds are less trustworthy due to the smaller amount of predicted MWEs.

6.3.3. Probabilistic identification

The previous section has considered the use of predicted compositionality scores as part of a rule-based system of MWE identification. We considered different thresholds on these scores, and we showed that higher thresholds produced better results. An analysis of the data suggested that higher scores may also be associated with higher rates of annotation (as they may indicate the presence of collocations). Rather than pursuing the rule-based approach under different ranges of threshold, we consider a different approach: using a probabilistic classifier which considers these scores as features for the prediction of MWE occurrences.

We construct a CRF classifier based on different kinds of features. The features, described in Section 6.2.4 can be purely contextual (CTX), or involve statistical association measures (AM), or come from the previously defined lexicon of compositionality scores (COMP_{W2V} and $\text{COMP}_{\text{PPMI}}$).

Table 6.7 presents the results obtained for the evaluation against the English reference corpus. Overall, there is a notable improvement in both precision and recall when additional features are considered beyond CTX. Concerning the recall, it can be seen that the use of association measures is redundant with the scores from COMP_{W2V} (as the recall in the second and third lines are identical, and the one in the fourth line is only slightly higher). The scores from $\text{COMP}_{\text{PPMI}}$, on the other hand, contribute to a higher improvement in the recall than the association measures. Moreover, this feature completely subsumes AM, as can be seen from the fact that the recall in the last two lines is the same. Regarding the precision, COMP scores are considerably higher than CTX alone, and to a certain extent CTX + AM as well. The highest F_1 score of .710 is better than the baseline .609 from Section 6.3.1 by +10.1 percentage points.

Features	Precision	Recall	F_1
CTX	.758	.567	.649
CTX + AM	.789	.602	.683
CTX + COMP_{W2V}	.797	.602	.686
CTX + COMP_{W2V} + AM	.798	.606	.689
CTX + $\text{COMP}_{\text{PPMI}}$.811	.626	.707
CTX + $\text{COMP}_{\text{PPMI}}$ + AM	.820	.626	.710

Table 6.7.: Results for probability-based MWE identifier (English dataset).

Table 6.8 presents the results obtained for the evaluation against the French reference corpus. As in the case of the English data above, the probabilistic method improves both precision and recall when we consider features beyond CTX. However, the improvement for the French corpus is much less pronounced. For recall, the greatest improvement happens with the addition of association measures (with +2.2 percentage points), with COMP providing a smaller effect.

In the case of the precision, only the addition of both AM and $\text{COMP}_{\text{PPMI}}$ provided an improvement, and it was considerably smaller than the one seen for English. Differently from what was observed for English, we see that the highest F_1 score obtained for the probabilistic method for French (.736) is unable to surpass the rule-based baseline of .763 obtained through $\text{TRAIN}_{\text{CONTIG}}$ in Section 6.3.1. Further work should investigate this behavior through an analysis of the French corpus annotations.

Features	Precision	Recall	F_1
CTX	.817	.636	.715
CTX + AM	.818	.658	.730
CTX + COMP_{W2V}	.802	.649	.717
CTX + COMP_{W2V} + AM	.812	.664	.731
CTX + $\text{COMP}_{\text{PPMI}}$.817	.650	.724
CTX + $\text{COMP}_{\text{PPMI}}$ + AM	.826	.664	.736

Table 6.8.: Results for probability-based MWE identifier (French dataset).

6.4. Summary

This chapter evaluated the use of automatically predicted compositionality scores as features in the task of MWE identification in two corpora. We started with a rule-based baseline, where a contiguous identification of lemmas seen in the training corpus was found to obtain high precision (higher than .8) for both languages. In the case of the English data, the recall from this method alone was weak (.23), while the recall obtained for French was considerably higher (.68), making this a particularly hard baseline to beat in the latter language. Indeed, while overall English results could be improved with the addition of a pattern for matching proper nouns (and the precision score could further be improved through noun–noun patterns), no F_1 improvement was found for the French corpus. The highest-scoring set of rules in this baseline obtained an $F_1 = .609$ for English and $F_1 = .763$ for French.

We then considered the application of compositionality scores directly as part of the rule-based method of MWE identification. We collected a lexicon of potential MWEs (based on NC syntactic patterns) and calculated their compositionality scores. Different thresholds were then applied on the compositionality scores, with the least-compositional MWEs being automatically annotated in the corpus according to the rule-based method. The results we obtained were consistent across the two languages and the two lexicons of compositionality: an overall low precision of identification, which surprisingly drops more harshly when stricter thresholds of idiomaticity are considered. We presume that this effect is

caused by a high rate of annotation of collocations, which tend to have higher compositionality scores.

Another method of identifying MWEs would be through a probabilistic approach considering multiple features. We evaluate the performance of a CRF trained on different sets of features, grouped as: lexical features, association measures, and compositionality scores. The results were highly corpus-dependent: while both association measures and compositionality scores contributed to higher values of F_1 for English, the improvement in French results was considerably weaker. Moreover, while the highest CRF scores for English (.710) convincingly beat the baseline above (by +10.1 percentage points), the highest scores for French (.736) are actually *lower* than the baseline of purely identifying all MWEs seen in training data (by -2.7 points). Further analysis would be required to understand this discrepancy.

Concerning the hypothesis $h_{\text{pred-comp}} \rightarrow \text{ident-accur}$ that compositionality scores can have a positive effect on the accuracy of MWE identification, we have obtained mixed results. In the case of rule-based methods, we found no improvement in prediction with more restrictive scores of compositionality (i.e. with a lower threshold). In fact, results suggest that the English corpus contains a high amount of annotated collocations, which would explain why a more strict threshold does not improve on the results. However, when we considered a CRF, the F_1 score for MWE identification for the English and French corpus did present an increase (in particular for the former language). In order to evaluate this hypothesis in a more favorable setting, future work should consider corpora that have been annotated particularly with idiomatic MWEs in mind. For example, the corpus for the PARSEME shared task contains an MWE category called *verbal idiom*, which is guaranteed to refer to expressions that humans have judged as idiosyncratic (Savary, Ramisch, Cordeiro, Sangati, et al. 2017), and could be a more close fit for the evaluation of this hypothesis.

One question that can be raised from the results in this chapter is whether association measures can be helpful in the identification of MWEs. While we did not find any correlation between human judgments of compositionality and a measure of conventionalization (see $h_{\text{idiom}} \approx \text{distr.convent}$ on page 61), note that the distinction being done on the task of compositionality prediction is between compositional and idiomatic MWEs, while the annotations on these corpora might tend toward a distinction between fully productive expressions and any kind of conventionalized expressions (i.e. any kind of MWEs, in the broadest sense). In this case, we hypothesize that the improvements in MWE identification caused by association measures is related to their ability of capturing conventionalization. Future analysis is still needed to verify whether this interpretation is correct. In particular, if only the idiomatic MWEs in the corpora are taken into account (i.e. compositional cases are filtered out), we do not expect association measures to contribute with an improvement of MWE detection scores.

The MWE identification techniques presented in this chapter were implemented

and are currently available as part of the mwetoolkit. A description of the implementation as well as further results can have been published as Cordeiro, Ramisch, and Villavicencio ([2016c](#)) and Scholivet, Ramisch, and Cordeiro ([2017](#)).

7. Conclusions

This thesis has proposed a framework of multiword expression compositionality prediction, and has investigated the impact of several variables in the accuracy of the predictions. The predictive model is based on the manipulation of distributional semantic models, i.e. vectorial representations of the meaning of words and MWEs. We have presented three new datasets of human-rated compositionality scores, in three different languages, and evaluate the developed framework using these resources. Finally, we also consider the use of predicted compositionality scores as features in the task of MWE identification.

Both the construction of the datasets and the subsequent evaluations of the predictive model are associated with a set of hypotheses. Table 7.1 summarizes these hypotheses and provides a reference to the page in which they have been evaluated.

In the following section, we present the main contributions from this thesis, including an overview of our findings for the evaluated hypotheses. We then present some perspectives of future work.

7.1. Contributions

The contributions of this thesis can be summarized as follows:

- Three new human-rated datasets of compositionality scores.
- An analysis of the new datasets with regards to score distribution and correlation with human variables.
- A new framework of compositionality prediction, which relies on a systematization of DSMs and parameters.
- A large-scale multilingual evaluation of the compositionality prediction framework under a variety of settings.
- An extrinsic evaluation of predicted compositionality scores in the task of MWE identification.

Many of the results presented in this thesis have also been presented in peer-reviewed publications. We refer back to Section 1.4 (page 21) for the complete list of publications. As for the contributions that have been presented in this thesis, we described them in more detail below.

Chapter 3 presented the construction of three datasets of human-rated MWE compositionality scores. The datasets encompass three languages (English, French

Hypothesis	Sub-hypothesis	Evaluation
$h_{\text{idiom}} \approx \text{distr}$	$h_{\text{idiom}} \approx \text{distr.convent}$	Pages 61, 106
	$h_{\text{idiom}} \approx \text{distr.freq}$	Pages 61, 103
$h_{\text{accur}} \leftarrow \text{MWE}$	$h_{\text{accur}} \leftarrow \text{MWE.diffic}$	Page 108
	$h_{\text{accur}} \leftarrow \text{MWE.idiom}$	Page 108
	$h_{\text{accur}} \leftarrow \text{MWE.convent}$	Page 106
	$h_{\text{accur}} \leftarrow \text{MWE.freq}$	Page 104
$h_{\text{accur}} \leftarrow \text{DSM}$	$h_{\text{accur}} \leftarrow \text{DSM.window}$	Page 81
	$h_{\text{accur}} \leftarrow \text{DSM.dims}$	Page 83
$h_{\text{accur}} \leftarrow \text{corpus}$	$h_{\text{accur}} \leftarrow \text{corpus.wordform}$	Page 84
	$h_{\text{accur}} \leftarrow \text{corpus.size}$	Page 86
	$h_{\text{accur}} \leftarrow \text{corpus.parallel}$	Page 87
h_{strat}	$h_{\text{strat.partial-info}}$	Page 90
	$h_{\text{strat.maxsim}}$	Page 91
	$h_{\text{strat.geom}}$	Page 94
$h_{\text{pred-comp}} \approx \text{comp}$	$h_{\text{pred-comp}} \approx \text{comp}$	Chapter 5
$h_{\text{pred-comp}} \rightarrow \text{ident-accur}$	$h_{\text{pred-comp}} \rightarrow \text{ident-accur}$	Chapter 6

Table 7.1.: Hypotheses evaluated in this thesis.

and Portuguese), and is the first dataset of MWE compositionality for two of these languages. This resource is freely available, and can be used for evaluating and training techniques that involve some type of semantic processing, such as lexical substitution and text simplification.

We also analyzed the constructed datasets, whose scores were found to follow a uniform distribution. Moreover, the three datasets were found to have comparable levels of difficulty of annotation. We have evaluated the hypothesis that MWE idiomaticity was correlated with distributional characteristics ($h_{\text{idiom}} \approx \text{distr}$). In particular, we considered the correlation with an estimator of conventionalization ($h_{\text{idiom}} \approx \text{distr.convent}$), which was shown not to be statistically significant. We also considered the correlation with the frequency ($h_{\text{idiom}} \approx \text{distr.freq}$), which turned out to be the opposite of what one would expect: higher-frequency MWEs are actually more likely to be compositional, with lower-frequency ones being more likely to be idiomatic. Some of these results were also presented in publications (Cordeiro, Ramisch, and Villavicencio 2016a; Wilkens, Zilio, Cordeiro, et al. 2017; Ramisch, Cordeiro, Zilio, et al. 2016).

Chapter 4 presents a systematization of different DSMs and their parameters along a common set of axes, which can be used to compare multiple distributional representations in a multilingual setting. This chapter also proposes a framework of compositionality prediction that can take into account all of these configurations. The framework has also been described in a publication (Cordeiro, Ramisch, and Villavicencio 2016b).

The framework above was then used in a large-scale intrinsic evaluation of multiple combinations of DSM and corpus parameters (Chapter 5). Here too we investigate the correlation between MWE idiomaticity and distributional characteristics ($h_{\text{idiom}} \approx \text{distr}$). As in the case of human judgments above, a correlation with corpus frequency was also observed for model predictions. Moreover, while human judgments did not correlate with a measure of conventionalization, we did find a mild correlation in some system scores.

One of the goals of the large-scale evaluation was to determine the factors that influence the accuracy of model predictions. One of the hypothesis we considered was that the accuracy should be influenced by MWE-specific characteristics ($h_{\text{accur}} \leftarrow \text{MWE}$). We found that both idiomaticity ($h_{\text{accur}} \leftarrow \text{MWE.idiom}$) and difficulty in human judgments of compositionality ($h_{\text{accur}} \leftarrow \text{MWE.diffic}$) are associated with lower-quality predictions. On the other hand, MWE frequency ($h_{\text{accur}} \leftarrow \text{MWE.freq}$) and conventionalization ($h_{\text{accur}} \leftarrow \text{MWE.convent}$) did not show clear signs of correlation with model accuracy.

We also evaluate different variations of DSMs. Our hypothesis is that DSM-specific configuration should play a crucial role in the accuracy of the results ($h_{\text{accur}} \leftarrow \text{DSM}$). While we do find a high variety in the accuracy across different DSMs, the results for the two DSM parameters we considered were somewhat underwhelming. We found that a higher number of dimensions would consistently contribute to a mild improvement in the accuracy ($h_{\text{accur}} \leftarrow \text{DSM.dims}$), but no cross-lingual and unified recommendations could be attained regarding the variation in context-window sizes ($h_{\text{accur}} \leftarrow \text{DSM.window}$).

Along with the impact from DSM-specific parameters, we also hypothesized an influence of corpus-specific parameters in the accuracy of results ($h_{\text{accur}} \leftarrow \text{corpus}$). Indeed, the results confirm that stopwords removal and lemmatization are both important steps of corpus preprocessing for this task, especially in the case of languages that are morphologically richer than English. Moreover, the use of POS tags does not contribute to a higher quality in the representation of word vectors for this task, possibly due to the fact that it increases the sparsity in co-occurrence counts ($h_{\text{accur}} \leftarrow \text{corpus.wordform}$). An analysis of different corpus sizes also showed that these may have a direct impact in the accuracy of results ($h_{\text{accur}} \leftarrow \text{corpus.size}$). Moreover, a proposed technique of parallel predictions was shown to perform equivalently to whole-corpus predictions, while allowing for the better utilization of computational resources ($h_{\text{accur}} \leftarrow \text{corpus.parallel}$).

Concerning the predictive model itself, we have considered six different strategies for deriving the compositionality scores. Our hypothesis is that different strategies would provide a different view into the data, with some strategies being more accurate than others (h_{strat}). We evaluated two additive strategies that are commonly used in the literature, but that had never been compared, and we concluded that their results are mostly equivalent to each other. Two other strategies considered only one of the words in the NCs (head or modifier). As expected, their accuracy suffered due to the limited information ($h_{\text{strat.partial-info}}$). We

then evaluated two proposed strategies. We confirmed the hypotheses that the *maxsim* strategy is better suited for compositional MWEs ($h_{\text{strat.maxsim}}$), while the *geom* strategy optimizes towards idiomatic cases ($h_{\text{strat.geom}}$). However, in both cases, results were quite similar to the standard additive strategies, suggesting that the impact of strategy choice is not as strong as previously thought.

All of the experiments presented in Chapter 5 revolve around a common hypothesis: model predictions are correlated with human-rated MWE compositionality ($h_{\text{pred-comp}} \approx \text{comp}$). Indeed, the variety of results obtained in this thesis all suggest that this hypothesis is true. While the correlation obtained for predictions using the worst DSM and corpus configurations may be considered weak, we have shown that the appropriate configurations are able to consistently produce predictions of compositionality that highly correlate with human judgments for a variety of datasets across multiple languages. The identification of patterns in the large space of more than 8 thousand configurations is one of the most salient contributions of this thesis. Some of the results on the evaluation of compositionality prediction were also published in an ACL paper (Cordeiro, Ramisch, Idiart, et al. 2016), and we are currently working on another paper to be submitted to a journal.

Finally, one of the contributions of this thesis is the application of predicted compositionality scores to the task of MWE identification (Chapter 6). The goal was to evaluate the hypothesis $h_{\text{pred-comp} \rightarrow \text{ident-accur}}$, which predicts an improvement of MWE identification with the use of predicted compositionality scores as internal features. In a rule-based algorithm, compositionality scores were not found to be a good feature for the identification of annotated MWE occurrences, likely due to the presence of collocations along with idiomatic MWEs in the annotation. We then considered a probabilistic model of identification, in which the results were mixed: while compositionality scores significantly improved the results over rule-based and probabilistic baselines for the English corpus, no such improvement was found for the French corpus. Further analysis of this phenomenon is left for future work. Intermediary results for this task have been published in Cordeiro, Ramisch, and Villavicencio (2016c), and sent for publication in Scholivet, Ramisch, and Cordeiro (2017).

7.2. Future work

Concerning the research on *compositionality datasets*, we envisage the extension of the dataset for each of the languages to allow better inter-language comparability (e.g. EN *red wine* and its translations FR *vin rouge* and PT *vinho tinto*). We also consider the collection of compositionality judgments for MWEs in additional languages, ideally from different language families for a broader generalization of results.

A different direction is to augment the dataset with judgments of similarity

between compounds sharing the same head. For example, we can ask people to judge the similarity of the word *case* in the expression *nut case* against the word *case* in similar expressions with high PMI, such as *criminal case*, *special case*, *exceptional case*, *upper case* and *business case*. This judgment could also be extended to all pairs of expressions, which would allow for semantic clusters (and where clusters of a single expression could be taken as evidence of idiomaticity). This approach could steer some of the research on compositionality in the direction of lexical similarity, which is commonly used for the evaluation of DSMs in the case of single words. It would also allow further investigation of polysemy in the case of collocations sharing the same head (Moldovan, Badulescu, Tatu, et al. 2004; Kim and Baldwin 2013). Crucially, it would allow us to peek into the DSM representation of similarly-looking compounds and to identify the ways in which the vectors of these expressions denote their difference in idiomaticity.

This thesis presents an extensive evaluation of an additive model of compositionality prediction using the constructed datasets. Similar evaluations could be done for other predictive techniques in the state of the art, such as the work of Salehi, Cook, and Baldwin (2015). The examination of these works in the context of our multilingual datasets would provide a more solid indication of their accuracy. Moreover, outstanding cross-language differences between such results and the results found in this thesis would provide further directions of investigation.

Regarding the results obtained in our *compositionality prediction* methods, the highest-scoring configurations in this thesis achieved reasonably high correlation with human predictions. Nevertheless, some of the predicted MWE scores were diametrically opposite to the average of human judgments. A cross-DSM analysis of vector representations for these MWEs could reveal whether this behavior stems from deficiencies in the underlying DSM vector representation (e.g. the fact that all DSMs considered could only represent a single meaning per word). If this is the case, modifications in the DSM could be investigated so as to prevent the occurrence of such discrepant score predictions. The construction of a dataset of compound head similarity such as the one suggested above could facilitate the discovery of these DSM weak points.

As for the compositionality prediction methods themselves, we plan on examining the use of a voting scheme for combining the output of complementary DSMs. Moreover, we also plan on combining additional sources of information for building the models, such as multilingual lexicons or translation data (Salehi, Cook, and Baldwin 2014), to improve even further the compositionality prediction. We would also like to propose and evaluate more sophisticated compositionality functions that take into account the unbalanced contribution of individual words to the global meaning of a compound. This could be done e.g. through a combination of the *maxsim* and *geom* strategies proposed in this thesis (either on the level of the strategy itself, or in the form of an ensemble method that combines the predictions of multiple strategies).

This thesis has employed predicted compositionality scores in an *application* of MWE identification. We considered a rule-based and a probabilistic model, both of which we evaluated under a base configuration as well as in two configurations involving compositionality scores. For the probabilistic model, technical considerations required the quantization of the predicted scores for a categorical interpretation. The specific quantization used may have greatly limited the results, and future works should consider different schemes of quantization. Alternatively, this problem could be solved through the use of neural networks, which can appropriately deal with real-valued data.

Considering the results obtained by the application of MWE identification, we see that compositionality scores significantly contribute to better accuracy in the case of the English corpus, but has a less pronounced effect on the French corpus. Crucially, the probabilistic method fares worse than the baseline which identifies only MWEs seen in the training data for the French language. Future work would be needed to investigate this difference between the results for the two languages.

Finally, we also consider other applications of compositionality scores. In particular, we would like to incorporate the collected scores into a machine translation system, as an indication of whether an expression should be translated as a single indivisible unit. We also envisage the application of predicted MWE compositionality scores in MWE-aware parsers, extending the approach used in previous work on multiword prepositions (Nasr, Ramisch, Deulofeu, et al. 2015; Constant and Nivre 2016; Waszczuk, Savary, and Parmentier 2016).

For the task of MWE identification, we would like to explore context-based definitions of compositionality scores. We would also like to evaluate our framework on verbal MWEs, such as the ones annotated for the PARSEME shared task (Savary, Ramisch, Cordeiro, Sangati, et al. 2017). Verbal MWEs are an understudied topic in the literature, and present some challenges that were not present in the case of the nominal compounds we used in this thesis. In particular, verbal MWEs can have extremely rigid or flexible morphosyntactic characteristics¹, and can often present discontinuities (e.g. *take [something] into account*). Work on verbal MWEs could be pursued in the context long-term research projects.

¹ For example, compare the rigid expression *bite me!*, which does not even allow the inflection of the verb, with the expression *pay a visit*, which even allows a change in word order in the passive voice.

Bibliography

- [ACT03] Anne Abeillé, Lionel Clément, and François Toussenet. “Building a treebank for French”. In: *Treebanks: building and using parsed corpora*. Ed. by Anne Abeillé. Dordrecht, The Netherlands: Kluwer academic publishers, 2003, pp. 165–168 (cit. on p. 118).
- [A+13] Shaishav Agrawal, Ankit Aggarwal, et al. “Hybrid Approach: A Solution for Extraction of Domain Independent Multiword Expression”. In: *Int J. of Technology Innovations and Research* 5 (2013) (cit. on p. 41).
- [AP08] Ron Artstein and Massimo Poesio. “Inter-Coder Agreement for Computational Linguistics”. In: *Computational Linguistics* 34.4 (2008), pp. 555–596. ISSN: 0891-2017 (cit. on p. 33).
- [BK10] Timothy Baldwin and Su Nam Kim. “Multiword Expressions”. In: *Handbook of Natural Language Processing, Second Edition*. 2010, pp. 267–292 (cit. on pp. 15, 32, 38, 39).
- [BV02] Timothy Baldwin and Aline Villavicencio. “Extracting the Unextractable: A Case Study on Verb-particles”. In: *Proceedings of CoNLL 2002. COLING-02. ACL, 2002*, pp. 1–7. DOI: [10.3115/1118853.1118854](https://doi.org/10.3115/1118853.1118854). URL: <http://dx.doi.org/10.3115/1118853.1118854> (cit. on p. 50).
- [Ban06] Colin James Bannard. “Acquiring phrasal lexicons from corpora”. PhD thesis. University of Edinburgh, 2006 (cit. on p. 50).
- [BBL03] Colin Bannard, Timothy Baldwin, and Alex Lascarides. “A Statistical Approach to the Semantics of Verb-particles”. In: *Proceedings of the ACL 2003 Workshop on Multiword Expressions: Analysis, Acquisition and Treatment - Volume 18. MWE '03*. Sapporo, Japan: Association for Computational Linguistics, 2003, pp. 65–72. DOI: [10.3115/1119282.1119291](https://doi.org/10.3115/1119282.1119291). URL: <http://dx.doi.org/10.3115/1119282.1119291> (cit. on p. 66).
- [Bar+09] Marco Baroni, Silvia Bernardini, Adriano Ferraresi, et al. “The WaCky wide web: a collection of very large linguistically processed web-crawled corpora”. In: *Language resources and evaluation* 43.3 (2009), pp. 209–226 (cit. on pp. 54, 71).
- [BDK14] Marco Baroni, Georgiana Dinu, and Germán Kruszewski. “Don’t count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors”. In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Baltimore, Maryland: Association for Computational

- Linguistics, June 2014, pp. 238–247. URL: <http://www.aclweb.org/anthology/P14-1023> (cit. on pp. 14, 46, 82).
- [BL10] Marco Baroni and Alessandro Lenci. “Distributional memory: A general framework for corpus-based semantics”. In: *Computational Linguistics* 36.4 (2010), pp. 673–721 (cit. on pp. 14, 15).
- [BZ10] Marco Baroni and Roberto Zamparelli. “Nouns are Vectors, Adjectives are Matrices: Representing Adjective-Noun Constructions in Semantic Space”. In: *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. Cambridge, MA: Association for Computational Linguistics, Oct. 2010, pp. 1183–1193. URL: <http://www.aclweb.org/anthology/D10-1115> (cit. on p. 43).
- [Bic00] Eckhard Bick. “The parsing system Palavras”. In: *Automatic Grammatical Analysis of Portuguese in a Constraint Grammar Framework* (2000) (cit. on p. 71).
- [BPV14] Rodrigo Boos, Kassius Prestes, and Aline Villavicencio. “Identification of Multiword Expressions in the brWaC”. In: *Proceedings of LREC 2014*. ACL Anthology Identifier: L14-1429. ELRA, May 2014, pp. 728–735. ISBN: 978-2-9517408-8-4. URL: http://www.lrec-conf.org/proceedings/lrec2014/pdf/518_Paper.pdf (cit. on pp. 54, 71).
- [Bot+16] Stefan Bott, Nana Khvtisavrishvili, Max Kisselew, et al. “GhoSt-PV: A Representative Gold Standard of German Particle Verbs”. In: *COLING 2016* (2016), p. 125 (cit. on p. 51).
- [Bou09] Gerlof Bouma. “Normalized (pointwise) mutual information in collocation extraction”. In: *Proceedings of GSCL* (2009), pp. 31–40 (cit. on p. 32).
- [BVA15] Antoine Bride, Tim Van de Cruys, and Nicholas Asher. “A Generalisation of Lexical Functions for Composition in Distributional Semantics.” In: *ACL (1)*. 2015, pp. 281–291 (cit. on p. 15).
- [BZL10] Fan Bu, Xiaoyan Zhu, and Ming Li. “Measuring the Non-compositionality of Multiword Expressions”. In: *Proceedings of the 23rd International Conference on Computational Linguistics*. COLING ’10. Beijing, China: Association for Computational Linguistics, 2010, pp. 116–124. URL: <http://dl.acm.org/citation.cfm?id=1873781.1873795> (cit. on pp. 18, 62).

- [BL12] John A Bullinaria and Joseph P Levy. “Extracting semantic representations from word co-occurrence statistics: stop-lists, stemming, and SVD”. In: *Behavior Research Methods* 44.3 (2012), pp. 890–907. ISSN: 1554-3528. DOI: [10.3758/s13428-011-0183-8](https://doi.org/10.3758/s13428-011-0183-8). URL: <http://dx.doi.org/10.3758/s13428-011-0183-8> (cit. on pp. 84, 85).
- [CPN15] José Camacho-Collados, Mohammad Taher Pilehvar, and Roberto Navigli. “A Framework for the Construction of Monolingual and Cross-lingual Word Similarity Datasets”. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Beijing, China: Association for Computational Linguistics, July 2015, pp. 1–7. URL: <http://www.aclweb.org/anthology/P15-2001> (cit. on p. 14).
- [Cap17] Fabienne Cap. “Show Me Your Variance and I Tell You Who You Are—Deriving Compound Compositionality from Word Alignments”. In: *MWE 2017* (2017), p. 102 (cit. on p. 67).
- [Cap+15] Fabienne Cap, Manju Nirmal, Marion Weller, et al. “How to Account for Idiomatic German Support Verb Constructions in Statistical Machine Translation”. In: *Proceedings of the 11th Workshop on Multiword Expressions*. Denver, Colorado: Association for Computational Linguistics, June 2015, pp. 19–28. URL: <http://www.aclweb.org/anthology/W15-0903> (cit. on pp. 15, 64).
- [CD10] Marine Carpuat and Mona Diab. “Task-based evaluation of multiword expressions: a pilot study in statistical machine translation”. In: *Proc. of NAACL/HLT 2010*. Los Angeles, CA, 2010, pp. 242–245 (cit. on p. 15).
- [CH90] Kenneth Ward Church and Patrick Hanks. “Word association norms, mutual information, and lexicography”. In: *Computational linguistics* 16.1 (1990), pp. 22–29 (cit. on p. 31).
- [CJ03] Massimiliano Ciaramita and Mark Johnson. “Supersense Tagging of Unknown Nouns in WordNet”. In: *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*. EMNLP ’03. Stroudsburg, PA, USA: Association for Computational Linguistics, 2003, pp. 168–175. DOI: [10.3115/1119355.1119377](https://doi.org/10.3115/1119355.1119377). URL: <https://doi.org/10.3115/1119355.1119377> (cit. on p. 43).
- [Con+17] Mathieu Constant, Gülşen Eryiğit, Johanna Monti, et al. “Multiword Expression Processing: A Survey”. In: *Computational Linguistics* (2017) (cit. on pp. 40–42, 115).
- [CN16] Matthieu Constant and Joakim Nivre. “A Transition-Based System for Joint Lexical and Syntactic Analysis.” In: *ACL*. 2016 (cit. on pp. 42, 135).

- [CS11] Matthieu Constant and Anthony Sigogne. “MWU-Aware Part-of-Speech Tagging with a CRF Model and Lexical Resources”. In: *Proceedings of the Workshop on Multiword Expressions: from Parsing and Generation to the Real World*. Portland, Oregon, USA: Association for Computational Linguistics, June 2011, pp. 49–56. URL: <http://www.aclweb.org/anthology/W11-0809> (cit. on p. 42).
- [Cor+16] Silvio Ricardo Cordeiro, Carlos Ramisch, Marco Idiart, et al. “Predicting the compositionality of nominal compounds: Giving word embeddings a hard time”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 2016, pp. 1986–1997 (cit. on pp. 17, 133).
- [CRV15] Silvio Ricardo Cordeiro, Carlos Ramisch, and Aline Villavicencio. “Token-based MWE Identification Strategies in the mwetoolkit”. In: *PARSEME’s 4th general meeting*. Malta, 2015 (cit. on p. 116).
- [CRV16a] Silvio Ricardo Cordeiro, Carlos Ramisch, and Aline Villavicencio. “Filtering and Measuring the Intrinsic Quality of Human Compositionality Judgments”. In: *Proceedings of the 12th Workshop on Multiword Expressions*. 2016, pp. 32–37 (cit. on pp. 17, 53, 63, 102, 131).
- [CRV16b] Silvio Ricardo Cordeiro, Carlos Ramisch, and Aline Villavicencio. “mwetoolkit+sem: Integrating Word Embeddings in the mwetoolkit for Semantic MWE Processing”. In: *Proc. of LREC 2016*. Portoroz, Slovenia, May 2016 (cit. on pp. 17, 71, 131).
- [CRV16c] Silvio Ricardo Cordeiro, Carlos Ramisch, and Aline Villavicencio. “UFRGS&LIF at SemEval-2016 Task 10: Rule-Based MWE Identification and Predominant-Supersense Tagging.” In: *SemEval at NAACL-HLT*. 2016, pp. 910–917 (cit. on pp. 17, 115, 120, 129, 133).
- [Dun93] Ted Dunning. “Accurate Methods for the Statistics of Surprise and Coincidence”. In: *Comput. Linguist.* 19.1 (Mar. 1993), pp. 61–74. ISSN: 0891-2017. URL: <http://dl.acm.org/citation.cfm?id=972450.972454> (cit. on p. 32).
- [EO15] Ismail El Maarouf and Michael Oakes. “Statistical Measures for Characterising MWEs”. In: *IC1207 COST PARSEME 5th general meeting*. 2015 (cit. on pp. 18, 62).
- [EP10] Katrin Erk and Sebastian Padó. “Exemplar-Based Models for Word Meaning in Context”. In: *ACL 2010, Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, July 11-16, 2010, Uppsala, Sweden, Short Papers*. The Association for Computer Linguistics, 2010, pp. 92–97. URL: <http://www.aclweb.org/anthology/P10-2017> (cit. on p. 14).

- [EK05] Stefan Evert and Brigitte Krenn. “Using small random samples for the manual evaluation of statistical association measures”. In: *Computer Speech & Language* 19.4 (2005), pp. 450–466 (cit. on p. 41).
- [FSN15] Meghdad Farahmand, Aaron Smith, and Joakim Nivre. “A Multiword Expression Data Set: Annotating Non-Compositionality and Conventionalization for English Noun Compounds”. In: *Proceedings of the 11th Workshop on Multiword Expressions*. Denver, Colorado: Association for Computational Linguistics, June 2015, pp. 29–33. URL: <http://www.aclweb.org/anthology/W15-0904> (cit. on pp. 32, 49, 51, 67).
- [FS06] Afsaneh Fazly and Suzanne Stevenson. “Automatically Constructing a Lexicon of Verb Phrase Idiomatic Combinations.” In: *EACL*. 2006 (cit. on pp. 18, 62).
- [Fel98] Christiane Fellbaum, ed. *WordNet: An Electronic Lexical Database (Language, Speech, and Communication)*. 423 p. MITPRESS, May 1998. ISBN: 0-262-06197-X (cit. on pp. 43, 48).
- [Fer14] Olivier Ferret. “Compounds and distributional thesauri”. In: *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*. Ed. by Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, et al. ACL Anthology Identifier: L14-1590. Reykjavik, Iceland: European Language Resources Association (ELRA), May 2014, pp. 2979–2984. ISBN: 978-2-9517408-8-4. URL: http://www.lrec-conf.org/proceedings/lrec2014/pdf/754_Paper.pdf (cit. on p. 15).
- [Fir57] John R Firth. “A synopsis of linguistic theory, 1930-1955”. In: (1957) (cit. on pp. 14, 44).
- [FAM00] Katerina Frantzi, Sophia Ananiadou, and Hideki Mima. “Automatic Recognition of Multi-Word Terms: the C-value/NC-value Method”. In: *International Journal on Digital Libraries* 3.2 (2000), pp. 115–130 (cit. on p. 41).
- [Fre60] Gottlob Frege. “Über Sinn und Bedeutung”. In: *Zeitschrift für Philosophie und philosophische Kritik* 100 (1892/1960). Translated, as ‘On Sense and Reference’, by Max Black, pp. 25–50 (cit. on p. 14).
- [Fre+05] Dayne Freitag, Matthias Blume, John Byrnes, et al. “New experiments in distributional representations of synonymy”. In: *Proceedings of the Ninth Conference on Computational Natural Language Learning*. Association for Computational Linguistics. 2005, pp. 25–32 (cit. on p. 14).

- [GVC13] Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. “PPDB: The Paraphrase Database.” In: *HLT-NAACL*. 2013, pp. 758–764 (cit. on p. 44).
- [Gir+05] Roxana Girju, Dan Moldovan, Marta Tatu, et al. “On the semantics of noun compounds”. In: *Computer speech & language* 19.4 (2005), pp. 479–496 (cit. on p. 48).
- [Gol15] Adele E. Goldberg. “Compositionality”. In: *The Routledge Handbook of Semantics*. Ed. by Nick Riemer. Routledge, 2015. Chap. 24 (cit. on p. 14).
- [Gue11] Emiliano Guevara. “Computing Semantic Compositionality in Distributional Semantics”. In: *Proceedings of the Ninth International Conference on Computational Semantics*. IWCS ’11. Oxford, United Kingdom: Association for Computational Linguistics, 2011, pp. 135–144. URL: <http://dl.acm.org/citation.cfm?id=2002669.2002684> (cit. on p. 66).
- [GA13] Antton Gurrutxaga and Iñaki Alegria. “Combining Different Features of Idiomaticity for the Automatic Classification of Noun+Verb Expressions in Basque”. In: *Proceedings of the 9th Workshop on Multiword Expressions*. Atlanta, Georgia, USA: Association for Computational Linguistics, June 2013, pp. 116–125. URL: <http://www.aclweb.org/anthology/W13-1017> (cit. on pp. 18, 51, 62).
- [Har54] Zellig Harris. “Distributional structure”. In: *Word* 10 (2-3 1954), pp. 146–162. DOI: [10.1080/00437956.1954.11659520](https://doi.org/10.1080/00437956.1954.11659520) (cit. on p. 14).
- [Har+17] Matthias Hartung, Fabian Kaupmann, Soufian Jebbara, et al. “Learning Compositionality Functions on Word Embeddings for Modelling Attribute Meaning in Adjective-Noun Phrases”. In: *Proceedings of the 15th Meeting of the European Chapter of the Association for Computational Linguistics (EACL)*. 2017 (cit. on p. 66).
- [Hen+13] Iris Hendrickx, Zornitsa Kozareva, Preslav Nakov, et al. “SemEval-2013 Task 4: Free Paraphrases of Noun Compounds”. In: *Proceedings of *SEM 2013, Volume 2 – SemEval*. ACL, June 2013, pp. 138–143. URL: <http://www.aclweb.org/anthology/S13-2025> (cit. on p. 48).
- [HXY15] Zhiheng Huang, Wei Xu, and Kai Yu. “Bidirectional LSTM-CRF models for sequence tagging”. In: *arXiv preprint arXiv:1508.01991* (2015) (cit. on p. 121).
- [Hwa+10] Jena D Hwang, Archana Bhatia, Clare Bonial, et al. “Propbank annotation of multilingual light verb constructions”. In: *Proc. of the LAW 2010*. ACL. 2010, pp. 82–90 (cit. on p. 15).

- [Jac97] Ray Jackendoff. *The Architecture of the Language Faculty*. Ed. by S. J. Keyser. Vol. 28. Cambridge, MA: MIT Press, 1997, pp. xvi+262+. ISBN: 0-262-60025-0 (cit. on pp. 15, 39).
- [JP15] Glorianna Jagfeld and Lonneke van der Plas. “Towards a better Semantic Role Labelling of Complex Predicates”. In: *Proc. of NAACL Student Research Workshop*. Denver, US, 2015, pp. 33–39 (cit. on p. 15).
- [JM08] Daniel Jurafsky and James H. Martin. *Speech and Language Processing*. 2nd. 1024 p. Upper Saddle River, NJ, USA: Prentice Hall, 2008. ISBN: 0-13-187321-0 (cit. on pp. 25, 27, 32).
- [JK95] John S. Justeson and Slava M. Katz. “Technical terminology: some linguistic properties and an algorithm for identification in text”. In: *Natural language engineering* 1.01 (1995), pp. 9–27 (cit. on p. 41).
- [KC14] Douwe Kiela and Stephen Clark. “A systematic study of semantic vector space model parameters”. In: *Proceedings of the 2nd Workshop on Continuous Vector Space Models and their Compositionality (CVSC) at EACL*. 2014, pp. 21–30 (cit. on pp. 74, 81, 84).
- [KB13] Su Nam Kim and Timothy Baldwin. “Word Sense and Semantic Relations in Noun Compounds”. In: *ACM Trans. Speech Lang. Process.* 10.3 (July 2013), 9:1–9:17. ISSN: 1550-4875. DOI: [10.1145/2483969.2483971](https://doi.org/10.1145/2483969.2483971). URL: <http://doi.acm.org/10.1145/2483969.2483971> (cit. on p. 134).
- [KW16] Maximilian Köper and Sabine Schulte im Walde. “Distinguishing Literal and Non-Literal Usage of German Particle Verbs.” In: *HLT-NAACL*. 2016, pp. 353–362 (cit. on p. 67).
- [KF11] Nidhi Kulkarni and Mark Finlayson. “jMWE: A Java Toolkit for Detecting Multi-Word Expressions”. In: *Proceedings of the Workshop on Multiword Expressions: From Parsing and Generation to the Real World*. MWE ’11. Portland, Oregon: Association for Computational Linguistics, 2011, pp. 122–124 (cit. on pp. 15, 42).
- [LMP01] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. “Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data”. In: *Proceedings of the Eighteenth International Conference on Machine Learning (ICML 2001)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001, pp. 282–289. ISBN: 1-55860-778-1. URL: <http://dl.acm.org/citation.cfm?id=645530.655813> (cit. on p. 117).
- [LFL98] Thomas K Landauer, Peter W Foltz, and Darrell Laham. “An introduction to latent semantic analysis”. In: *Discourse processes* 25.2-3 (1998), pp. 259–284 (cit. on p. 14).

- [LLT98] Shari Landes, Claudia Leacock, and Randee I Tengi. “Building semantic concordances”. In: *WordNet: an electronic lexical database* 199.216 (1998), pp. 199–216 (cit. on pp. 43, 48).
- [LE14] Gabriella Lapesa and Stefan Evert. “A Large Scale Evaluation of Distributional Semantic Models: Parameters, Interactions and Model Selection”. In: *Transactions of the Association for Computational Linguistics* 2 (2014), pp. 531–545 (cit. on pp. 74, 75, 81, 82).
- [LE17] Gabriella Lapesa and Stefan Evert. “Large-scale evaluation of dependency-based DSMs: Are they worth the effort?” In: *EACL 2017* (2017), p. 394 (cit. on pp. 15, 82).
- [Lau95] Mark Lauer. “How much is enough?: Data requirements for statistical NLP”. In: *CoRR* abs/cmp-lg/9509001 (1995). URL: <http://arxiv.org/abs/cmp-lg/9509001> (cit. on p. 48).
- [LGD15] Omer Levy, Yoav Goldberg, and Ido Dagan. “Improving Distributional Similarity with Lessons Learned from Word Embeddings”. In: *Transactions of the Association for Computational Linguistics* 3 (2015), pp. 211–225. ISSN: 2307-387X. URL: <https://tacl2013.cs.columbia.edu/ojs/index.php/tacl/article/view/570> (cit. on pp. 46, 73, 84, 97).
- [Lin98] Dekang Lin. “Automatic retrieval and clustering of similar words”. In: *Proceedings of the 17th international conference on Computational linguistics-Volume 2*. Association for Computational Linguistics. 1998, pp. 768–774 (cit. on pp. 14, 46).
- [Lin99] Dekang Lin. “Automatic identification of non-compositional phrases”. In: *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*. Association for Computational Linguistics. 1999, pp. 317–324 (cit. on p. 14).
- [MS99] Christopher D. Manning and Hinrich Schütze. *Foundations of statistical natural language processing*. 620 p. Cambridge, USA, 1999. ISBN: 0-262-13360-1 (cit. on p. 27).
- [MKC03] Diana McCarthy, Bill Keller, and John Carroll. “Detecting a Continuum of Compositionality in Phrasal Verbs”. In: *Proceedings of the ACL 2003 Workshop on Multiword Expressions*. MWE ’03. Sapporo, Japan: ACL, 2003, pp. 73–80. DOI: [10.3115/1119282.1119292](https://doi.org/10.3115/1119282.1119292). URL: <http://dx.doi.org/10.3115/1119282.1119292> (cit. on pp. 15, 50, 67).
- [Mik+13] Tomas Mikolov, Ilya Sutskever, Kai Chen, et al. “Distributed representations of words and phrases and their compositionality”. In: *Advances in neural information processing systems*. 2013, pp. 3111–3119 (cit. on pp. 14, 15, 19, 47, 73, 84, 109).

- [ML08] Jeff Mitchell and Mirella Lapata. “Vector-based Models of Semantic Composition.” In: *ACL*. 2008, pp. 236–244 (cit. on pp. 66, 67).
- [ML10] Jeff Mitchell and Mirella Lapata. “Composition in distributional models of semantics”. In: *Cognitive science* 34.8 (2010), pp. 1388–1429 (cit. on pp. 15, 19, 68, 109).
- [Mit05] Ruslan Mitkov. *The Oxford handbook of computational linguistics*. Oxford University Press, 2005 (cit. on p. 27).
- [Mol+04] Dan Moldovan, Adriana Badulescu, Marta Tatu, et al. “Models for the Semantic Classification of Noun Phrases”. In: *Proceedings of the HLT-NAACL Workshop on Computational Lexical Semantics*. CLS ’04. Boston, Massachusetts: Association for Computational Linguistics, 2004, pp. 60–67. URL: <http://dl.acm.org/citation.cfm?id=1596431.1596440> (cit. on p. 134).
- [Nak08] Preslav Nakov. “Paraphrasing verbs for noun compound interpretation”. In: *Proc. of the LREC Workshop Towards a Shared Task for MWEs (MWE 2008)*. 2008, pp. 46–49 (cit. on p. 48).
- [Nak13] Preslav Nakov. “On the interpretation of noun compounds: Syntax, semantics, and entailment”. In: *Natural Language Engineering* 19.3 (2013), pp. 291–330. URL: <http://dx.doi.org/10.1017/S1351324913000065> (cit. on p. 40).
- [Nas+15] Alexis Nasr, Carlos Ramisch, José Deulofeu, et al. “Joint Dependency Parsing and Multiword Expression Tokenisation”. In: *Annual Meeting of the Association for Computational Linguistics*. 2015, pp. 1116–1126 (cit. on p. 135).
- [NP12] Roberto Navigli and Simone Paolo Ponzetto. “BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network”. In: *Artificial Intelligence* 193 (2012), pp. 217–250 (cit. on p. 44).
- [NHN06] Joakim Nivre, Johan Hall, and Jens Nilsson. “Maltparser: A data-driven parser-generator for dependency parsing”. In: *Proceedings of LREC*. Vol. 6. 2006, pp. 2216–2219 (cit. on p. 71).
- [Oka07] Naoaki Okazaki. *CRFsuite: a fast implementation of Conditional Random Fields (CRFs)*. 2007. URL: <http://www.chokkan.org/software/crfsuite/> (cit. on p. 117).
- [PL07] Sebastian Padó and Mirella Lapata. “Dependency-based construction of semantic space models”. In: *Computational Linguistics* 33.2 (2007), pp. 161–199 (cit. on p. 46).

- [Pad+14] Muntsa Padró, Marco Idiart, Aline Villavicencio, et al. “Nothing like Good Old Frequency: Studying Context Filters for Distributional Thesauri”. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP 2014) - short papers*. Doha, Qatar, Oct. 2014. URL: <http://www.aclweb.org/anthology/D14-1047> (cit. on pp. 46, 73).
- [Par14] Barbara H Partee. *Montague grammar*. Elsevier, 2014 (cit. on p. 43).
- [Pec10] Pavel Pecina. “Lexical association measures and collocation extraction”. In: *Language resources and evaluation* 44.1-2 (2010), pp. 137–158 (cit. on p. 32).
- [PSM14] Jeffrey Pennington, Richard Socher, and Christopher Manning. “Glove: Global Vectors for Word Representation”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics, Oct. 2014, pp. 1532–1543. URL: <http://www.aclweb.org/anthology/D14-1162> (cit. on pp. 46, 73, 84, 97).
- [PDM11] Slav Petrov, Dipanjan Das, and Ryan McDonald. “A universal part-of-speech tagset”. In: *arXiv preprint arXiv:1104.2086* (2011) (cit. on p. 25).
- [Pin95] Steven Pinker. “Language acquisition”. In: *Language: An invitation to cognitive science* 1 (1995), pp. 135–82 (cit. on pp. 18, 61).
- [Ram15] Carlos Ramisch. *Multiword Expressions Acquisition - A Generic and Open Framework*. Theory and Applications of Natural Language Processing. Springer, 2015. DOI: 10.1007/978-3-319-09207-2. URL: <http://dx.doi.org/10.1007/978-3-319-09207-2> (cit. on pp. 21, 41, 54).
- [Ram+16] Carlos Ramisch, Silvio Ricardo Cordeiro, Leonardo Zilio, et al. “How Naked is the Naked Truth? A Multilingual Lexicon of Nominal Compound Compositionality”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 2016, p. 156 (cit. on pp. 17, 53, 63, 131).
- [RM95] Lance Ramshaw and Mitch Marcus. “Text Chunking using Transformation-Based Learning”. In: *Third Workshop on Very Large Corpora*. 1995. URL: <http://aclweb.org/anthology/W95-0107> (cit. on p. 118).
- [RMM11] Siva Reddy, Diana McCarthy, and Suresh Manandhar. “An Empirical Study on Compositionality in Compound Nouns”. In: *Proceedings of The 5th International Joint Conference on Natural Language Processing 2011 (IJCNLP 2011)*. Chiang Mai, Thailand, Nov. 2011. URL: <http://sivareddy.in/papers/ijcnlp2011empirical.pdf> (cit. on pp. 15, 16, 49, 50, 53, 55–57, 66, 69, 102, 113, 155).

- [Ren+09] Zhixiang Ren, Yajuan Lü, Jie Cao, et al. “Improving Statistical Machine Translation Using Domain Bilingual Multiword Expressions”. In: *Proc. of the ACL 2009 Workshop on MWEs*. Singapore, 2009, pp. 47–54 (cit. on p. 15).
- [RN11] Denise Rey and Markus Neuhäuser. “Wilcoxon-signed-rank test”. In: *International encyclopedia of statistical science*. Springer, 2011, pp. 1658–1659 (cit. on p. 76).
- [RB15] Martin Riedl and Chris Biemann. “A Single Word is not Enough: Ranking Multiword Expressions Using Distributional Semantics.” In: *EMNLP*. 2015, pp. 2430–2440 (cit. on p. 67).
- [RS14] Stephen Roller and Sabine Schulte im Walde. “Feature Norms of German Noun Compounds”. In: *Proceedings of the 10th Workshop on Multiword Expressions (MWE)*. ACL, Apr. 2014, pp. 104–108. URL: <http://www.aclweb.org/anthology/W14-0818> (cit. on p. 51).
- [RSS13] Stephen Roller, Sabine Schulte im Walde, and Silke Scheible. “The (Un)expected Effects of Applying Standard Cleansing Models to Human Ratings on Compositionality”. In: *Proceedings of the 9th Workshop on Multiword Expressions*. ACL, June 2013, pp. 32–41. URL: <http://www.aclweb.org/anthology/W13-1005> (cit. on pp. 51, 100, 102).
- [Sag+02] Ivan A Sag, Timothy Baldwin, Francis Bond, et al. “Multiword expressions: A pain in the neck for NLP”. In: *Computational Linguistics and Intelligent Text Processing*. Springer, 2002, pp. 1–15 (cit. on pp. 15, 38, 39).
- [SCB14] Bahar Salehi, Paul Cook, and Timothy Baldwin. “Using Distributional Similarity of Multi-way Translations to Predict Multiword Expression Compositionality”. In: *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*. Ed. by Gosse Bouma and Yannick Parmentier. Gothenburg, Sweden: The Association for Computer Linguistics, Apr. 2014, pp. 472–481. URL: <http://aclweb.org/anthology/E/E14/E14-1050.pdf> (cit. on pp. 46, 67, 113, 134).
- [SCB15] Bahar Salehi, Paul Cook, and Timothy Baldwin. “A Word Embedding Approach to Predicting the Compositionality of Multiword Expressions”. In: *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Denver, Colorado: Association for Computational Linguistics, May 2015, pp. 977–983. URL: <http://www.aclweb.org/anthology/N15-1099> (cit. on pp. 16, 66, 70, 73, 76, 111, 113, 134).

- [Sal+15] Bahar Salehi, Nitika Mathur, Paul Cook, et al. “The Impact of Multiword Expression Compositionality on Machine Translation Evaluation”. In: *Proceedings of the 11th Workshop on Multiword Expressions*. Denver, Colorado: Association for Computational Linguistics, June 2015, pp. 54–59. URL: <http://www.aclweb.org/anthology/W15-0909> (cit. on pp. 15, 64).
- [SVI16] Alexandre Salle, Aline Villavicencio, and Marco Idiart. “Matrix Factorization using Window Sampling and Negative Sampling for Improved Word Representations”. In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 419–424. URL: <http://anthology.aclweb.org/P16-2068> (cit. on pp. 46, 73, 84, 97).
- [Sav09] Agata Savary. “Multiflex: A Multilingual Finite-State Tool for Multi-Word Units.” In: *CIAA*. Ed. by Sebastian Maneth. Vol. 5642. Lecture Notes in Computer Science. Springer, 2009, pp. 237–240. ISBN: 978-3-642-02978-3. URL: <http://dblp.uni-trier.de/db/conf/wia/ciaa2009.html#Savary09> (cit. on p. 42).
- [Sav+17a] Agata Savary, Marie Candito, Verginica Barbu Mititelu, et al. “PARSEME multilingual corpus of verbal multiword expressions”. In: *Phraseology and Multiword Expressions*. Submitted for review. Language Science Press (LangSci), 2017.
- [Sav+17b] Agata Savary, Carlos Ramisch, Silvio Ricardo Cordeiro, Marie Candito, et al. “The PARSEME shared task on automatic identification of verbal multiword expressions”. In: *Phraseology and Multiword Expressions*. Submitted for review. Language Science Press (LangSci), 2017 (cit. on p. 42).
- [Sav+17c] Agata Savary, Carlos Ramisch, Silvio Ricardo Cordeiro, Federico Sangati, et al. “The PARSEME Shared Task on Automatic Identification of Verbal Multiword Expressions”. In: *Proceedings of the 13th Workshop on Multiword Expressions (MWE 2017)*. 2017, pp. 31–47 (cit. on pp. 21, 128, 135).
- [Sav+15] Agata Savary, Manfred Sailer, Yannick Parmentier, et al. “PARSEME – PARSing and Multiword Expressions within a European multilingual network”. In: *7th Language & Technology Conference: Human Language Technologies as a Challenge for Computer Science and Linguistics (LTC 2015)*. Poznań, Poland, Nov. 2015 (cit. on p. 40).
- [Sch95] Helmut Schmid. “Treectagger — a language independent part-of-speech tagger”. In: *Institut für Maschinelle Sprachverarbeitung, Universität Stuttgart 43* (1995), p. 28 (cit. on p. 71).

- [Sch14] Nathan Schneider. “Lexical Semantic Analysis in Natural Language Text”. PhD thesis. Ph. D. dissertation, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, 2014 (cit. on p. 118).
- [Sch+14a] Nathan Schneider, Emily Danchik, Chris Dyer, et al. “Discriminative lexical semantic segmentation with gaps: running the MWE gamut”. In: *Transactions of the Association for Computational Linguistics* 2 (Apr. 2014), pp. 193–206. URL: <http://www.transacl.org/wp-content/uploads/2014/04/51.pdf> (cit. on p. 42).
- [Sch+16a] Nathan Schneider, Dirk Hovy, Anders Johannsen, et al. “SemEval 2016 Task 10: Detecting Minimal Semantic Units and their Meanings (DiMSUM)”. In: *Proc. of SemEval*. San Diego, California, USA, June 2016 (cit. on pp. 15, 40, 42, 49, 115, 118).
- [Sch+16b] Nathan Schneider, Jena D Hwang, Vivek Srikumar, et al. “A corpus of preposition supersenses”. In: *Proceedings of the 10th Linguistic Annotation Workshop held in conjunction with ACL*. 2016, pp. 99–109 (cit. on pp. 43, 49).
- [Sch+14b] Nathan Schneider, Spencer Onuffer, Nora Kazour, et al. “Comprehensive Annotation of Multiword Expressions in a Social Web Corpus”. In: *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*. Ed. by Nicoletta Calzolari, Khalid Choukri, Thierry Declerck, et al. ACL Anthology Identifier: L14-1433. Reykjavik, Iceland: European Language Resources Association (ELRA), May 2014, pp. 455–461. ISBN: 978-2-9517408-8-4. URL: http://www.lrec-conf.org/proceedings/lrec2014/pdf/521_Paper.pdf (cit. on pp. 115, 118).
- [Sch+14c] Nathan Schneider, Spencer Onuffer, Nora Kazour, et al. “Comprehensive annotation of multiword expressions in a social web corpus”. In: (2014) (cit. on p. 49).
- [SRC17] Manon Scholivet, Carlos Ramisch, and Silvio Ricardo Cordeiro. “Sequence Models and Lexical Resources for MWE Identification in French”. In: *Phraseology and Multiword Expressions*. Submitted for review. Language Science Press (LangSci), 2017 (cit. on pp. 17, 42, 115, 121, 129, 133).
- [SJ01] Patrick Schone and Daniel Jurafsky. “Is Knowledge-Free Induction of Multiword Unit Dictionary Headwords a Solved Problem?” In: *Proceedings of Empirical Methods in Natural Language Processing*. Pittsburgh, PA, 2001. URL: http://www.colorado.edu/ling/jurafsky/emnlp2001_mwu_iii.pdf (cit. on p. 67).
- [Sch+16c] Sabine Schulte im Walde, Anna Hättig, Stefan Bott, et al. “GhoSt-NN: A Representative Gold Standard of German Noun-Noun Compounds.” In: *LREC*. 2016 (cit. on p. 51).

- [SMR13] Sabine Schulte im Walde, Stefan Müller, and Stefan Roller. “Exploring Vector Space Models to Predict the Compositionality of German Noun-Noun Compounds”. In: *Proceedings of *SEM 2013, Volume 1*. ACL, 2013, pp. 255–265. URL: <http://www.aclweb.org/anthology/S13-1038> (cit. on pp. 66, 70).
- [Sed+13] Djamé Seddah, Reut Tsarfaty, Sandra Kübler, et al. “Overview of the SPMRL 2013 Shared Task: A Cross-Framework Evaluation of Parsing Morphologically Rich Languages”. In: *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages*. Seattle, WA, USA: Association for Computational Linguistics, 2013, pp. 146–182. URL: <http://www.aclweb.org/anthology/W13-4917> (cit. on p. 118).
- [Ser11] Violeta Seretan. *Syntax-Based Collocation Extraction*. 1st. Vol. 44. Text, Speech and Language Technology. 212 p. Dordrecht, Netherlands: Springer, 2011. ISBN: 978-94-007-0133-5 (cit. on p. 41).
- [Shl14] Jonathon Shlens. “A tutorial on principal component analysis”. In: *arXiv preprint arXiv:1404.1100* (2014) (cit. on p. 46).
- [Sma93] Frank Smadja. “Retrieving Collocations from Text: Xtract”. In: *Comput. Linguist.* 19.1 (Mar. 1993), pp. 143–177. ISSN: 0891-2017. URL: <http://dl.acm.org/citation.cfm?id=972450.972458> (cit. on p. 41).
- [Soc+12] Richard Socher, Brody Huval, Christopher D Manning, et al. “Semantic compositionality through recursive matrix-vector spaces”. In: *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. Association for Computational Linguistics. 2012, pp. 1201–1211 (cit. on p. 15).
- [SL09] Caroline Sporleder and Linlin Li. “Unsupervised Recognition of Literal and Non-Literal Use of Idiomatic Expressions”. In: *Proc. of EACL 2009*. Athens, 2009, pp. 754–762 (cit. on p. 67).
- [SFN04] Suzanne Stevenson, Afsaneh Fazly, and Ryan North. “Statistical Measures of the Semi-productivity of Light Verb Constructions”. In: *Proceedings of the Workshop on Multiword Expressions: Integrating Processing*. MWE '04. Barcelona, Spain: Association for Computational Linguistics, 2004, pp. 1–8. URL: <http://dl.acm.org/citation.cfm?id=1613186.1613187> (cit. on p. 41).
- [SCA13] Sara Stymne, Nicola Cancedda, and Lars Ahrenberg. “Generation of compound words in statistical machine translation into compounding languages”. In: *Computational Linguistics* 39.4 (2013), pp. 1067–1108 (cit. on p. 64).

- [TW14] Yulia Tsvetkov and Shuly Wintner. “Identification of multiword expressions by combining multiple linguistic information sources”. In: *Computational Linguistics* 40.2 (2014), pp. 449–468 (cit. on p. 41).
- [WSP16] Jakub Waszczuk, Agata Savary, and Yannick Parmentier. “Promoting multiword expressions in A* TAG parsing”. In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. Osaka, Japan: The COLING 2016 Organizing Committee, Dec. 2016, pp. 429–439. URL: <http://aclweb.org/anthology/C16-1042> (cit. on p. 135).
- [Wil+17] Rodrigo Wilkens, Leonardo Zilio, Silvio Ricardo Cordeiro, et al. “Lex-SubNC: A Dataset of Lexical Substitution for Nominal Compounds”. In: *Proceedings of the Twelfth International Conference on Computational Semantics*. IWCS’17. Association for Computational Linguistics, 2017 (cit. on pp. 53, 56, 63, 131).
- [YFH15] Majid Yazdani, Meghdad Farahmand, and James Henderson. “Learning Semantic Composition to Detect Non-compositionality of Multiword Expressions”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, Sept. 2015, pp. 1733–1742. URL: <http://aclweb.org/anthology/D15-1201> (cit. on pp. 51, 66, 70, 75, 79, 113, 114).
- [Zho12] Zhi-Hua Zhou. *Ensemble methods: foundations and algorithms*. CRC press, 2012 (cit. on p. 87).
- [Zil+16] Leonardo Zilio, Rodrigo Wilkens, Luís Möllmann, et al. “Joining Forces for Multiword Expression Identification”. In: *International Conference on Computational Processing of the Portuguese Language*. Springer. 2016, pp. 233–238.

Index

Symbols

α , 33
 β , 66
 BF_1 , 37
 $c_H/c_M/c_{WC}$, 56
 \cos , 45
 CS_A , 69
 CS_β , 68
 CS_G , 70
 CS_P , 88
 D , 44
 δ , 73
 ϕ , 118
 F_1 , 37
 freq , 29
 κ , 33
 ρ , 35
 r , 35
 rank , 35
 rk , 92
 $\sigma_H/\sigma_M/\sigma_{WC}$, 56
 T , 117
 τ , 120
 V , 27
 v , 44
 w , 66
 y , 118
 $\#\text{occur}$, 28
 $\#\text{co-occur}$, 30

A

ALL-comp, 57
annotation, 33
association measure, 31

B

BF_1 score, 37
BIO, 118

C

CBOW, 47

co-occurrence, [29](#)
composition strategies, [68](#)
compound, [40](#)
content word, [28](#)
context word, [30](#)
conventionalization, [49](#)
corpus, [27](#), [71](#), [118](#)
correlation measure, [34](#)
cosine, [45](#)
covariance, [34](#)

D

dependent, [26](#)
dimension, [74](#)
DiMSUM, [40](#)

E

EN-comp₉₀, [53](#)

F

F₁ score, [37](#)
fallback evaluation, [76](#)
Farahmand, [51](#)
FR-comp, [53](#)
frequency, [29](#)
function word, [28](#)

G

GloVe, [73](#)
guidelines, [33](#)

H

head, [26](#)
hypotheses, [18](#)

I

improvement score, [92](#)

L

lemma, [25](#)
lexvec, [73](#)

M

multiword expression, [38](#)
MWE, [38](#)

N

NC, [40](#)

NLP, [14](#)

nominal compound, [40](#)

noun phrase, [26](#)

O

occurrence, [28](#)

P

Pearson's r , [35](#)

phrase, [26](#)

PMI, [31](#)

POS tag, [25](#)

PPMI, [32](#), [72](#)

precision, [37](#)

PT-comp, [53](#)

R

ranking, [35](#)

recall, [37](#)

Reddy, [51](#)

Reddy⁺⁺, [53](#)

S

Spearman's ρ , [35](#)

stopword, [28](#)

strategies, [68](#)

strict evaluation, [76](#)

supervised approach, [27](#)

surface form, [25](#)

SVD, [46](#)

T

target word, [30](#)

token, [25](#), [27](#)

tokenization, [25](#)

type, [27](#)

U

unsupervised approach, [27](#)

V

variance, [34](#)

vocabulary, [27](#)

W

WaC, [71](#)

window, [74](#)

word embeddings, [46](#)

word form, [74](#)

word2vec, [73](#)

APPENDIX

A. List of English Compounds

We present below the 90 nominal compounds in *EN-comp*₉₀, along with their human-rated compositionality scores. We refer to Reddy, McCarthy, and Manandhar (2011) for the other 90 compounds.

Compounds	cWC	Compounds	cWC
ancient history	1.95	high life	1.67
armchair critic	1.33	inner circle	1.56
baby buggy	3.94	inner product	3.00
bad hat	0.62	insane asylum	3.95
benign tumour	4.69	insurance company	5.00
big fish	0.85	insurance policy	4.15
birth rate	4.60	iron collar	3.88
black cherry	3.11	labour union	4.76
bow tie	4.25	life belt	2.84
brain teaser	2.65	life vest	3.44
busy bee	0.88	lime tree	4.61
carpet bombing	1.24	loan shark	1.00
cellular phone	3.78	loose woman	2.53
close call	1.59	mail service	4.69
closed book	0.68	market place	3.00
computer program	4.50	mental disorder	4.89
con artist	2.10	middle school	3.84
cooking stove	4.68	milk tooth	1.43
cotton candy	1.79	mother tongue	0.59
critical review	4.06	narrow escape	1.75
dead end	1.32	net income	2.94
dirty money	2.21	news agency	4.39
dirty word	2.48	noble gas	1.18
disc jockey	1.25	nut case	0.44
divine service	3.11	old flame	0.58
dry land	3.95	old hat	0.35
dry wall	3.33	old timer	0.89
dust storm	3.85	phone book	4.25
eager beaver	0.36	pillow slip	3.70
economic aid	4.33	pocket book	1.42
elbow grease	0.56	prison guard	4.89
elbow room	0.61	prison term	4.79
entrance hall	4.17	private eye	0.82
eternal rest	3.25	record book	3.70
fish story	1.68	research lab	4.75
flower child	0.50	sex bomb	0.53
food market	3.82	silver lining	0.35
foot soldier	1.95	sound judgement	3.39
front man	1.64	sparkling water	3.14
goose egg	0.48	street girl	3.16
grey matter	2.39	subway system	4.63
guinea pig	0.45	tennis elbow	2.50
half sister	2.84	top dog	1.05
half wit	1.16	wet blanket	0.21
health check	4.17	word painting	1.62

B. List of French Compounds

We present below the 180 nominal compounds in *FR-comp*, along with their human-rated compositionality scores.

Compounds	cwc	Gloss
activité physique	4.93	‘physical activity’ (lit. <i>activity physical</i>)
année scolaire	3.60	‘school year’ (lit. <i>year scholar</i>)
art contemporain	4.60	‘contemporary art’ (lit. <i>art contemporary</i>)
baie vitrée	3.64	‘open glass window’ (lit. <i>opening glassy</i>)
bas côté	1.31	‘aisle’ (lit. <i>low side</i>)
beau frère	0.67	‘brother-in-law’ (lit. <i>beautiful brother</i>)
beau père	1.18	‘father-in-law’ (lit. <i>beautiful father</i>)
belle mère	0.80	‘mother-in-law’ (lit. <i>beautiful mother</i>)
berger allemand	1.29	‘German shepherd’ (lit. <i>shepherd German</i>)
bon sens	3.57	‘common sense’ (lit. <i>good sense</i>)
bon vent	0.87	‘good luck’ (lit. <i>good/fair wind</i>)
bon vivant	2.57	‘bon vivant’ (lit. <i>good "liver"</i>)
bonne humeur	4.53	‘good mood’ (lit. <i>good humor</i>)
bonne poire	0.42	‘sucker, soft touch’ (lit. <i>good pear</i>)
bonne pratique	4.47	‘good practice’ (lit. <i>good practice</i>)
bouc émissaire	0.23	‘scapegoat’ (lit. <i>goat emissary</i>)
bras cassé	0.57	‘lame duck’ (lit. <i>arm broken</i>)
bras droit	0.40	‘right arm’ (lit. <i>arm right</i>)
brebis galeuse	0.55	‘black sheep’ (lit. <i>sheep scabby</i>)
carte blanche	0.20	‘carte blanche’ (lit. <i>card white</i>)
carte bleue	1.94	‘bank card’ (lit. <i>card blue</i>)
carte grise	3.08	‘vehicle registration’ (lit. <i>card grey</i>)
carte vitale	1.70	‘healthcare card’ (lit. <i>card vital</i>)
carton plein	0.78	‘clean sweep’ (lit. <i>cardboard full</i>)
casque bleu	1.85	‘UN peacekeeper’ (lit. <i>helmet blue</i>)
centre commercial	3.93	‘shopping center’ (lit. <i>center commercial</i>)
cercle vicieux	2.15	‘vicious circle’ (lit. <i>circle vicious</i>)
cerf volant	0.64	‘kite’ (lit. <i>deer flying</i>)
chambre froide	4.27	‘cold chamber’ (lit. <i>chamber cold</i>)
changement climatique	4.79	‘climate change’ (lit. <i>change climatic</i>)
chapeau bas	0.64	‘bravo’ (lit. <i>hat low</i>)
charge sociale	3.00	‘social security contribution’ (lit. <i>charge social</i>)
chauve souris	0.33	‘bat’ (lit. <i>bald mouse</i>)
chute libre	3.64	‘free fall’ (lit. <i>fall free</i>)
club privé	4.58	‘private club’ (lit. <i>club private</i>)
coffre fort	3.67	‘safe, vault’ (lit. <i>chest/box strong</i>)
communauté urbaine	4.57	‘urban community’ (lit. <i>community urban</i>)
conseil municipal	4.00	‘city council’ (lit. <i>council municipal</i>)
coup dur	2.40	‘hard blow’ (lit. <i>blow hard</i>)
coup franc	1.71	‘free kick (soccer)’ (lit. <i>blow free/frank</i>)
courrier électronique	4.57	‘e-mail’ (lit. <i>mail electronic</i>)
court circuit	1.69	‘short circuit’ (lit. <i>short circuit</i>)
court métrage	2.36	‘short film’ (lit. <i>short length/footage</i>)
crème fraîche	3.73	‘French sour cream’ (lit. <i>cream fresh</i>)
crème glacée	4.75	‘ice cream’ (lit. <i>cream icy</i>)
dernier cri	0.67	‘latest, trendy’ (lit. <i>last cry</i>)
dernier mot	3.09	‘final say’ (lit. <i>last word</i>)
directeur général	3.87	‘chief executive officer’ (lit. <i>director general</i>)
disque dur	2.83	‘hard drive’ (lit. <i>disk hard</i>)
douche froide	1.18	‘damper’ (lit. <i>cold shower</i>)
droit fondamental	4.27	‘fundamental right’ (lit. <i>right fundamental</i>)
développement économique	4.46	‘economic development’ (lit. <i>development economic</i>)
eau chaude	5.00	‘hot water’ (lit. <i>water hot</i>)
eau douce	2.33	‘fresh water’ (lit. <i>water sweet</i>)
eau minérale	4.00	‘mineral water’ (lit. <i>water mineral</i>)

Compounds	cWC	Gloss
eau potable	5.00	'drinking water' (lit. <i>water potable</i>)
eau vive	3.44	'jellyfish' (lit. <i>water living</i>)
eau forte	0.90	'etching' (lit. <i>water strong</i>)
eaux usées	4.54	'sewage' (lit. <i>waters used</i>)
effet spécial	3.67	'special effect' (lit. <i>effect special</i>)
expérience professionnelle	4.86	'professional experience' (lit. <i>experience professional</i>)
fait divers	3.69	'news story' (lit. <i>fact diverse</i>)
famille nombreuse	4.90	'large family' (lit. <i>family numerous</i>)
faux ami	1.25	'false friend' (lit. <i>false friend</i>)
faux cul	0.31	'hypocrite' (lit. <i>false arse</i>)
faux pas	1.82	'blunder' (lit. <i>false step</i>)
faux semblant	3.57	'false pretence' (lit. <i>false appearance</i>)
feu rouge	2.60	'red traffic light' (lit. <i>fire red</i>)
feu vert	0.71	'green light, permission' (lit. <i>fire green</i>)
fil conducteur	1.25	'underlying theme' (lit. <i>thread conducting</i>)
fleur bleue	0.45	'sentimental' (lit. <i>flower blue</i>)
foie gras	4.54	'foie gras' (lit. <i>liver fatty</i>)
fou rire	2.33	'giggle' (lit. <i>crazy laughter</i>)
grand air	1.33	'outdoors' (lit. <i>big air</i>)
grand jour	1.07	'broad daylight' (lit. <i>big day</i>)
grand saut	2.17	'move forward' (lit. <i>big leap</i>)
grand écran	3.14	'silver screen' (lit. <i>big screen</i>)
grande entreprise	4.54	'big company' (lit. <i>big company</i>)
grande surface	3.14	'department store' (lit. <i>big surface</i>)
grippe aviaire	3.58	'avian flu' (lit. <i>flu avian</i>)
gros mot	1.40	'swearword' (lit. <i>large word</i>)
gros plan	1.87	'close-up' (lit. <i>large plan</i>)
guerre civile	3.43	'civil war' (lit. <i>war civil</i>)
haut parleur	1.83	'loudspeaker' (lit. <i>loud/high speaker</i>)
haute mer	2.54	'high seas' (lit. <i>high sea</i>)
haute montagne	4.13	'high mountains' (lit. <i>high mountain</i>)
heure supplémentaire	4.00	'overtime hour' (lit. <i>hour extra</i>)
huile essentielle	2.25	'essential oil' (lit. <i>oil essential</i>)
idée reçue	2.90	'popular belief' (lit. <i>idea received</i>)
insertion professionnelle	4.27	'professional insertion' (lit. <i>insertion professional</i>)
intérêt général	4.36	'general interest' (lit. <i>interest general</i>)
jeune fille	4.64	'young girl, maiden' (lit. <i>young girl</i>)
journal officiel	4.50	'official gazette' (lit. <i>newspaper official</i>)
langue française	4.85	'French language' (lit. <i>language French</i>)
marée noire	3.00	'oil spill' (lit. <i>tide black</i>)
match nul	2.46	'draw, stalemate' (lit. <i>match null</i>)
matière grasse	5.00	'fat' (lit. <i>matter greasy</i>)
matière grise	2.15	'grey matter' (lit. <i>material grey</i>)
matière première	2.90	'raw material' (lit. <i>material primary</i>)
mauvaise foi	2.38	'bad faith' (lit. <i>bad faith</i>)
mauvaise langue	2.21	'gossip' (lit. <i>bad tongue</i>)
montagnes russes	1.08	'roller coaster' (lit. <i>mountains Russian</i>)
monument historique	4.79	'historical monument' (lit. <i>monument historical</i>)
mort né	3.23	'stillborn' (lit. <i>dead born</i>)
nouveau monde	2.73	'New World, Americas' (lit. <i>new world</i>)
nuit blanche	1.07	'sleepless night' (lit. <i>night white</i>)
numéro vert	1.50	'toll-free number' (lit. <i>number green</i>)
ordure ménagère	4.20	'household waste' (lit. <i>garbage household</i>)
organisation syndicale	4.90	'trade union' (lit. <i>organisation of-trade-union</i>)
pages jaunes	3.00	'yellow pages' (lit. <i>pages yellow</i>)
parachute doré	0.50	'golden parachute' (lit. <i>parachute golden</i>)
parc naturel	4.33	'nature park' (lit. <i>park natural</i>)
parti politique	4.88	'political party' (lit. <i>party political</i>)
parti pris	2.69	'bias' (lit. <i>party taken</i>)
partie fine	0.80	'orgy' (lit. <i>party fine/delicate</i>)
petit ami	0.86	'boyfriend' (lit. <i>small friend</i>)
petit beurre	1.64	'butter biscuit' (lit. <i>small butter</i>)

Compounds	cWC	Gloss
petit déjeuner	2.27	'breakfast' (lit. <i>small lunch</i>)
petit joueur	1.00	'amateur' (lit. <i>small player</i>)
petit pois	4.14	'pea' (lit. <i>small pea</i>)
petit salé	1.15	'salted pork' (lit. <i>small salty</i>)
petit écran	2.50	'television' (lit. <i>small screen</i>)
petit enfant	2.79	'grandchild' (lit. <i>small child</i>)
petit four	0.92	'type of dessert' (lit. <i>small oven</i>)
petit nègre	0.50	'pidgin French' (lit. <i>little black-man</i>)
petite annonce	2.69	'classified ad' (lit. <i>small announcement</i>)
petite nature	0.47	'squeamish' (lit. <i>small nature</i>)
piéd noir	0.13	'French expats from Algeria' (lit. <i>foot black</i>)
pièce montée	2.47	'tiered cake' (lit. <i>piece assembled</i>)
pleine lune	3.54	'full moon' (lit. <i>full moon</i>)
poids lourd	2.08	'truck' (lit. <i>weight heavy</i>)
point faible	2.46	'weak point' (lit. <i>point weak</i>)
point mort	1.00	'standstill' (lit. <i>point dead</i>)
pot pourri	0.40	'medley' (lit. <i>pot/jar rotten</i>)
poule mouillée	0.00	'coward' (lit. <i>chicken wet</i>)
poupée russe	3.75	'Russian nesting doll' (lit. <i>doll Russian</i>)
premier ministre	3.67	'first minister' (lit. <i>first minister</i>)
premier plan	2.82	'foreground' (lit. <i>first plan</i>)
première dame	1.92	'first lady' (lit. <i>first lady</i>)
prince charmant	2.00	'Prince Charming' (lit. <i>prince charming</i>)
prévision météorologique	4.70	'weather forecast' (lit. <i>forecast meteorological</i>)
recherche scientifique	4.92	'scientific research' (lit. <i>research scientific</i>)
ressources humaines	3.91	'human resources' (lit. <i>resources human</i>)
rond point	3.18	'roundabout' (lit. <i>round point</i>)
roulette russe	0.87	'Russian roulette' (lit. <i>roulette Russian</i>)
réchauffement climatique	4.40	'global warming' (lit. <i>warming climatic</i>)
région parisienne	4.43	'Paris region' (lit. <i>region Parisian</i>)
réseau social	4.09	'social network' (lit. <i>network social</i>)
sang froid	0.47	'self-control' (lit. <i>blood cold</i>)
second degré	1.40	'tongue-in-cheek' (lit. <i>second degree</i>)
second rôle	3.64	'supporting role' (lit. <i>second role</i>)
septième ciel	0.21	'cloud nine' (lit. <i>seventh heaven</i>)
service public	4.71	'public service' (lit. <i>service public</i>)
site officiel	4.85	'official website' (lit. <i>website official</i>)
soirée privée	4.53	'private party' (lit. <i>party private</i>)
sucre roux	4.31	'brown sugar' (lit. <i>sugar ginger-colored</i>)
sécurité routière	4.55	'road safety' (lit. <i>safety of-road</i>)
sécurité sociale	3.67	'social security' (lit. <i>security social</i>)
table basse	4.79	'coffee table' (lit. <i>table low</i>)
table ronde	1.46	'round table' (lit. <i>table round</i>)
tapis rouge	3.31	'red carpet' (lit. <i>carpet red</i>)
temps fort	1.87	'key moment, highlight' (lit. <i>time strong</i>)
temps mort	2.07	'wasted time, idleness' (lit. <i>time dead</i>)
temps partiel	3.62	'part-time (work)' (lit. <i>time partial</i>)
temps plein	3.08	'full-time (work)' (lit. <i>time full</i>)
temps réel	3.00	'real time' (lit. <i>time real</i>)
travaux publics	4.09	'public works' (lit. <i>works public</i>)
trou noir	2.58	'black hole' (lit. <i>hole black</i>)
trou normand	0.78	'palate cleanser' (lit. <i>hole Norman</i>)
téléphone arabe	0.23	'Chinese whispers' (lit. <i>telephone Arabic</i>)
téléphone portable	5.00	'cellphone' (lit. <i>telephone portable</i>)
valeur sûre	3.64	'safe bet' (lit. <i>value safe/sure</i>)
vie associative	4.00	'community life' (lit. <i>life associative</i>)
vie quotidienne	4.31	'everyday life' (lit. <i>life daily</i>)
vieille fille	2.42	'spinster' (lit. <i>old girl/maid</i>)
vin blanc	3.80	'white wine' (lit. <i>wine white</i>)
vin rouge	4.69	'red wine' (lit. <i>wine red</i>)
yeux rouges	4.36	'red eyes' (lit. <i>eyes red</i>)
école primaire	3.92	'primary school' (lit. <i>school primary</i>)
étoile filante	3.20	'shooting star' (lit. <i>star slipping</i>)

C. List of Portuguese Compounds

We present below the 180 nominal compounds in *PT-comp*, along with their human-rated compositionality scores.

Compounds	cwc	Gloss
abalo sísmico	4.42	‘earthquake’ (lit. <i>shock seismic</i>)
acampamento militar	4.82	‘military camp’ (lit. <i>camp military</i>)
agente secreto	4.58	‘secret agent’ (lit. <i>agent secret</i>)
alarme falso	3.24	‘false alarm’ (lit. <i>alarm false</i>)
algodão doce	1.28	‘cotton candy’ (lit. <i>cotton sweet</i>)
alta temporada	2.04	‘high season’ (lit. <i>high season</i>)
alta costura	1.52	‘haute couture’ (lit. <i>high sewing</i>)
alto mar	1.35	‘high seas’ (lit. <i>high sea</i>)
alto falante	0.88	‘loudspeaker’ (lit. <i>loud/high speaker</i>)
amigo oculto	2.89	‘secret Santa’ (lit. <i>friend hidden</i>)
amigo secreto	3.11	‘secret Santa’ (lit. <i>friend secret</i>)
amor próprio	3.91	‘self-esteem’ (lit. <i>love own</i>)
ano novo	4.29	‘new year’ (lit. <i>year new</i>)
ar condicionado	2.44	‘air conditioning’ (lit. <i>air conditioned</i>)
ar livre	1.95	‘open air’ (lit. <i>air free</i>)
arma branca	0.65	‘cold weapon’ (lit. <i>weapon white</i>)
ato falho	3.50	‘Freudian slip’ (lit. <i>act faulty</i>)
banho turco	2.19	‘Turkish bath’ (lit. <i>bath Turkish</i>)
batata doce	4.24	‘sweet potato’ (lit. <i>potato sweet</i>)
bebida alcoólica	5.00	‘alcoholic drink’ (lit. <i>drink alcoholic</i>)
bode expiatório	0.47	‘scapegoat’ (lit. <i>goat expiatory</i>)
braço direito	0.57	‘right arm’ (lit. <i>arm right</i>)
buraco negro	2.88	‘black hole’ (lit. <i>hole black/dark</i>)
café colonial	2.70	‘afternoon tea’ (lit. <i>breakfast colonial</i>)
caixa forte	3.19	‘safe, vault’ (lit. <i>box strong</i>)
caixa preta	0.94	‘black box’ (lit. <i>box black</i>)
caixeiro viajante	3.43	‘traveling salesman’ (lit. <i>clerk traveling</i>)
carne branca	2.85	‘white meat’ (lit. <i>meat white</i>)
carne vermelha	3.66	‘red meat’ (lit. <i>meat red</i>)
carro forte	2.62	‘armored car’ (lit. <i>car strong</i>)
carta aberta	3.64	‘open letter’ (lit. <i>letter open</i>)
centro comercial	3.68	‘shopping mall’ (lit. <i>center commercial</i>)
centro espírita	3.43	‘Spiritualist center’ (lit. <i>center spiritualist</i>)
cerca viva	3.58	‘hedge’ (lit. <i>fence living</i>)
cheiro verde	0.67	‘parsley’ (lit. <i>smell green</i>)
circuito integrado	4.52	‘integrated circuit’ (lit. <i>circuit integrated</i>)
classe executiva	2.67	‘business class’ (lit. <i>class executive</i>)
coluna social	2.45	‘gossip column’ (lit. <i>column social</i>)
colégio militar	4.88	‘military high-school’ (lit. <i>high-school military</i>)
comida caseira	4.11	‘homemade food’ (lit. <i>food homemade</i>)
companhia aérea	3.11	‘airline’ (lit. <i>company aerial</i>)
conta corrente	2.71	‘checking account’ (lit. <i>account current</i>)
coração partido	1.06	‘broken heart’ (lit. <i>heart broken</i>)
corda bamba	1.31	‘tightrope, bad situation’ (lit. <i>rope wobbly</i>)
cordas vocais	2.32	‘vocal chords’ (lit. <i>chords vocal</i>)
curto circuito	1.96	‘short circuit’ (lit. <i>short circuit</i>)
câmara fria	4.65	‘cold chamber’ (lit. <i>chamber cold</i>)
céu aberto	1.68	‘outdoors, open air’ (lit. <i>sky open</i>)
círculo vicioso	2.17	‘vicious circle’ (lit. <i>circle vicious</i>)
círculo virtuoso	2.39	‘virtuous circle’ (lit. <i>circle virtuous</i>)
deputado federal	4.92	‘federal deputy’ (lit. <i>deputy federal</i>)
desfile militar	4.93	‘military parade’ (lit. <i>parade military</i>)
direitos humanos	3.86	‘human rights’ (lit. <i>rights human</i>)
disco rígido	2.76	‘hard drive’ (lit. <i>disk rigid</i>)

Compounds	cWC	Gloss
disco voador	2.94	'flying saucer' (lit. <i>disk flying</i>)
efeitos especiais	3.37	'special effects' (lit. <i>effects special</i>)
elefante branco	0.16	'white elephant' (lit. <i>elephant white</i>)
escada rolante	3.85	'escalator' (lit. <i>stair rolling</i>)
estrela cadente	2.52	'shooting star' (lit. <i>star falling</i>)
exame clínico	4.75	'clinical examination' (lit. <i>examination clinical</i>)
exames laboratoriais	4.90	'laboratory tests' (lit. <i>examinations laboratory</i>)
farinha integral	4.72	'wholemeal flour' (lit. <i>flour integral</i>)
febre amarela	1.43	'yellow fever' (lit. <i>fever yellow</i>)
ficha limpa	2.97	'clean criminal records' (lit. <i>file clean</i>)
fila indiana	1.17	'single file' (lit. <i>queue Indian</i>)
fio condutor	1.58	'underlying theme' (lit. <i>thread conductor</i>)
força bruta	3.33	'brute force' (lit. <i>force brute</i>)
gatos pingados	0.00	'a few people' (lit. <i>cats dropped</i>)
gelo seco	2.33	'dry ice' (lit. <i>ice dry</i>)
golpe baixo	2.03	'low blow' (lit. <i>punch low</i>)
governo federal	4.97	'federal government' (lit. <i>government federal</i>)
gripe aviária	3.11	'avian flu' (lit. <i>flu avian</i>)
gripe suína	2.48	'swine flu' (lit. <i>flu swine</i>)
guarda florestal	4.16	'forest ranger' (lit. <i>guard forest</i>)
jogo duro	1.13	'rough play' (lit. <i>game hard</i>)
juízo final	3.60	'doomsday' (lit. <i>judgement final</i>)
leite integral	4.67	'whole milk' (lit. <i>milk integral</i>)
lista negra	1.60	'black list' (lit. <i>list black</i>)
livre-docente	2.63	'professor' (lit. <i>free lecturer</i>)
livro aberto	0.79	'open book' (lit. <i>book open</i>)
longa data	1.63	'longtime' (lit. <i>date long</i>)
longa-metragem	0.96	'feature film' (lit. <i>long length/footage</i>)
lua cheia	3.52	'full moon' (lit. <i>moon full</i>)
lua nova	1.40	'new moon' (lit. <i>moon new</i>)
lugar comum	1.52	'cliché' (lit. <i>place common</i>)
magia negra	1.72	'black magic' (lit. <i>magic black</i>)
mar aberto	2.87	'open sea' (lit. <i>sea open</i>)
maré alta	4.03	'high tide' (lit. <i>tide high</i>)
maré baixa	4.18	'low tide' (lit. <i>tide low</i>)
massa cinzenta	1.69	'grey matter' (lit. <i>mass grey</i>)
mau contato	2.84	'faulty contact' (lit. <i>bad contact</i>)
mau humor	4.29	'bad mood' (lit. <i>bad humour</i>)
mau olhado	1.97	'evil eye' (lit. <i>bad glance</i>)
mercado negro	1.06	'black market' (lit. <i>black market</i>)
mesa redonda	1.10	'round table' (lit. <i>table round</i>)
montanha russa	0.31	'roller coaster' (lit. <i>mountain Russian</i>)
má fé	1.62	'bad faith' (lit. <i>bad faith</i>)
máquina virtual	3.76	'virtual machine' (lit. <i>machine virtual</i>)
mão fechada	1.06	'stingy' (lit. <i>hand closed</i>)
navio negreiro	3.52	'slave ship' (lit. <i>ship black-slave</i>)
novo mundo	2.29	'new world' (lit. <i>new world</i>)
novo rico	3.62	'new rich, new money' (lit. <i>new rich</i>)
nó cego	0.74	'difficult situation' (lit. <i>knot blind</i>)
núcleo atômico	4.93	'atomic nucleus' (lit. <i>nucleus atomic</i>)
olho gordo	0.28	'evil eye' (lit. <i>eye fat</i>)
olho mágico	0.27	'peephole' (lit. <i>eye magic</i>)
olho nu	2.15	'naked eye' (lit. <i>eye naked</i>)
ovelha negra	0.45	'black sheep' (lit. <i>sheep black</i>)
papel higiênico	4.27	'toilet paper' (lit. <i>paper hygienic</i>)
paraíso fiscal	1.47	'tax haven' (lit. <i>paradise fiscal</i>)
pastor alemão	0.90	'German shepherd' (lit. <i>shepherd German</i>)
pau mandado	0.30	'subservient, stooge' (lit. <i>stick ordered</i>)
pavio curto	0.80	'short-tempered' (lit. <i>fuse short</i>)
pente fino	0.53	'careful research' (lit. <i>comb thin</i>)
peso morto	0.90	'dead weight' (lit. <i>weight dead</i>)
planta baixa	0.74	'floor plan' (lit. <i>plant short</i>)
ponto cego	1.92	'blind spot' (lit. <i>point blind</i>)

Compounds	cwc	Gloss
ponto forte	1.51	‘strong point’ (lit. <i>point strong</i>)
ponto fraco	2.27	‘weak point’ (lit. <i>point weak</i>)
poção mágica	3.29	‘magic potion’ (lit. <i>potion magic</i>)
prato feito	3.14	‘blue-plate special’ (lit. <i>plate ready-made</i>)
primeira infância	3.70	‘early childhood’ (lit. <i>first infancy</i>)
primeira-mão	0.71	‘first hand’ (lit. <i>first hand</i>)
primeira necessidade	3.97	‘first necessity’ (lit. <i>first necessity</i>)
primeira-dama	1.52	‘first lady’ (lit. <i>first dame</i>)
primeiro-ministro	2.87	‘first minister’ (lit. <i>first minister</i>)
primeiro plano	2.00	‘forefront’ (lit. <i>first plan</i>)
processo seletivo	4.78	‘selection process’ (lit. <i>process selective</i>)
pronto socorro	2.76	‘first-aid posts’ (lit. <i>ready aid</i>)
príncipe encantado	1.72	‘prince charming’ (lit. <i>prince enchanted</i>)
puro sangue	1.55	‘pure blood’ (lit. <i>pure blood</i>)
pão-duro	0.12	‘stingy’ (lit. <i>bread hard</i>)
pé quente	0.09	‘lucky’ (lit. <i>foot hot</i>)
pé-direito	0.10	‘ceiling height’ (lit. <i>foot right</i>)
pé frio	0.23	‘unlucky’ (lit. <i>foot cold</i>)
pólo aquático	2.87	‘water polo’ (lit. <i>aquatic pole/polo</i>)
quadro negro	2.94	‘blackboard’ (lit. <i>board black</i>)
queda livre	3.48	‘free fall’ (lit. <i>fall free</i>)
quinta categoria	1.00	‘second-rate’ (lit. <i>fifth category</i>)
rede social	3.27	‘social network’ (lit. <i>network social</i>)
regime político	4.00	‘political system’ (lit. <i>regime political</i>)
relógio analógico	4.92	‘analog clock’ (lit. <i>clock analog</i>)
relógio biológico	2.12	‘biological clock’ (lit. <i>clock biological</i>)
reta final	1.12	‘final stretch’ (lit. <i>straight line final</i>)
roda gigante	4.20	‘Ferris wheel’ (lit. <i>wheel giant</i>)
roleta russa	0.29	‘Russian roulette’ (lit. <i>roulette Russian</i>)
saia justa	0.37	‘tight spot’ (lit. <i>skirt tight</i>)
sala cirúrgica	4.47	‘operating room’ (lit. <i>room surgical</i>)
salão paroquial	4.52	‘parish hall’ (lit. <i>hall parish</i>)
sangue azul	0.15	‘blue-blooded’ (lit. <i>blood blue</i>)
sangue frio	0.52	‘cold-blooded’ (lit. <i>blood cold</i>)
sangue quente	0.87	‘hot-blooded’ (lit. <i>blood hot</i>)
secretária eletrônica	2.52	‘answering machine’ (lit. <i>secretary electronic</i>)
segundas intenções	2.11	‘ulterior motives’ (lit. <i>second intentions</i>)
segundo plano	1.55	‘aside, in the background’ (lit. <i>second plan</i>)
sentença judicial	4.67	‘court ruling’ (lit. <i>sentence judicial</i>)
sexto sentido	1.40	‘sixth sense’ (lit. <i>sixth sense</i>)
sinal verde	1.39	‘green lights’ (lit. <i>signal green</i>)
sistema político	4.36	‘political system’ (lit. <i>system political</i>)
sétima arte	2.19	‘seventh art’ (lit. <i>seventh art</i>)
tapete vermelho	3.76	‘red carpet’ (lit. <i>carpet red</i>)
tartaruga marinha	5.00	‘sea turtle’ (lit. <i>turtle marine</i>)
tela plana	4.96	‘flat screen TV’ (lit. <i>screen flat</i>)
tempo real	2.81	‘real time’ (lit. <i>time real</i>)
terceira idade	1.70	‘elder’ (lit. <i>third age</i>)
terceira pessoa	2.00	‘third person’ (lit. <i>third person</i>)
tiro livre	1.58	‘free kick (soccer)’ (lit. <i>shot free</i>)
trabalho braçal	3.55	‘manual labor’ (lit. <i>work arm</i>)
trabalho escravo	4.24	‘slave work’ (lit. <i>work slave</i>)
vaca louca	1.23	‘mad cow’ (lit. <i>cow crazy/mad</i>)
vinho branco	3.40	‘white wine’ (lit. <i>wine white</i>)
vinho tinto	4.08	‘red wine’ (lit. <i>wine dark-red</i>)
vista grossa	0.50	‘turn a blind eye’ (lit. <i>vision thick</i>)
viva voz	1.70	‘aloud’ (lit. <i>live voice</i>)
voto secreto	4.82	‘secret ballot’ (lit. <i>vote secret</i>)
vôo doméstico	3.41	‘domestic flight’ (lit. <i>flight domestic</i>)
vôo internacional	4.96	‘international flight’ (lit. <i>flight international</i>)
água doce	1.45	‘fresh water’ (lit. <i>water sweet</i>)
água mineral	4.21	‘mineral water’ (lit. <i>water mineral</i>)
ônibus executivo	2.63	‘minibus’ (lit. <i>bus executive</i>)