

Cross-Domain Ambiguity Detection using Linear Transformation of Word Embedding Spaces

Vaibhav Jain
vaibhav29498@gmail.com

Ruchika Malhotra
ruchikamalhotra@dtu.ac.in

Sanskar Jain
sanskar27jain@gmail.com

Nishant Tanwar
nishant.tanwar08@gmail.com

Department of Software Engineering
Delhi Technological University
India

Abstract

The requirements engineering process is a crucial stage of the software development life cycle. It involves various stakeholders from different professional backgrounds, particularly in the requirements elicitation phase. Each stakeholder carries distinct domain knowledge, causing them to differently interpret certain words, leading to cross-domain ambiguity. This can result in misunderstanding amongst them and jeopardize the entire project. This paper proposes a natural language processing approach to find potentially ambiguous words for a given set of domains. The idea is to apply linear transformations on word embedding models trained on different domain corpora, to bring them into a *unified embedding space*. The approach then finds words with divergent embeddings as they signify a variation in the meaning across the domains. It can help a requirements analyst in preventing misunderstandings during elicitation interviews and meetings by defining a set of potentially ambiguous terms in advance. The paper also discusses certain problems with the existing approaches and discusses how the proposed approach resolves them.

1 Introduction

In the context of software engineering, requirements engineering (RE) is the process of describing the intended behaviour of a software system along with the associated constraints [Pre10]. One of its phase is requirements elicitation, which has been termed as the most difficult, critical, and communication-intensive aspect of software development [AS05]. It requires interaction between different stakeholders through various techniques like brainstorming sessions and facilitated application specification technique. A stakeholder is any person with a vested interest in the project, such as potential users, developers, testers, domain experts, and regulatory agency personnel [SM12]. As these stakeholders come from different professional backgrounds and carry different domain

Copyright © 2020 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

In: M. Sabetzadeh, A. Vogelsang, S. Abualhaija, M. Borg, F. Dalpiaz, M. Daneva, N. Fernández, X. Franch, D. Fucci, V. Gervasi, E. Groen, R. Guizzardi, A. Herrmann, J. Horkoff, L. Mich, A. Perini, A. Susi (eds.): Joint Proceedings of REFSQ-2020 Workshops, Doctoral Symposium, Live Studies Track, and Poster Track, Pisa, Italy, 24-03-2020, published at <http://ceur-ws.org>

knowledge, cross-domain ambiguity can occur amongst them. One may assign an interpretation to another’s expression different from the intended meaning. This results in misunderstanding and distrust in requirements elicitation meetings, and costly problems in the later stages of the software life cycle [WMGWF13].

The study of variation of word meanings across domains as an NLP problem is termed as *Synchronic Lexical Semantic Change (LSC) Detection* [SHDTSIW19]. The first attempt to apply it for dealing with cross-domain ambiguity in RE was by Ferrari *et al.* (2017) who used Wikipedia crawling and word embeddings to estimate ambiguous computer science (CS) terms vis-à-vis other application domains [FDG17]. Mishra and Sharma extended this work by focusing on various engineering subdomains [MS19]. Another approach was suggested by Ferrari *et al.* (2018) which also considered the ambiguity caused by non-CS domain-specific words and addressed some of the technical limitations of the previous work [FEG18]. This approach was later extended to include quantitative evaluation of the obtained results [FE19]. An alternative approach which doesn’t require domain-specific word embeddings was suggested by Toews and Holland [TH19].

This paper proposes a natural language processing (NLP) approach based on linear transformation of word embedding spaces. Word embedding is a vector representation of a word capable of capturing its semantic and syntactic relations. A linear transformation can be used to learn a linear relationship between two word embedding spaces. The proposed approach produces a ranked list of potentially ambiguous terms for a given set of domains. It constructs a word embedding space for each domain using corpora composed of Wikipedia articles. It then applies linear transformations on these spaces in order to align them and construct a *unified embedding space*. For each word in a set of target words, an ambiguity score is assigned by applying a distance metric on its *domain-specific embeddings*.

The remainder of this paper is organised as follows: Section 2 provides some background on ambiguity in RE and linear transformation of word embedding spaces. The existing approaches to cross-domain ambiguity detection are briefly explained in Section 3. The motivation behind the proposed approach is discussed in Section 4, whereas the approach itself is outlined in Section 5. The experimental setup and results are presented and discussed in Section 6, and the conclusion and details of planned future work are provided in Section 7.

2 Preliminaries

2.1 Ambiguity in Requirements Engineering

Ambiguity refers to the ability of a natural language (NL) expression to be interpreted in multiple manners. As requirements elicitation is a communication-intensive process, ambiguity is a major negative factor as it can lead to an unclear and incomplete requirements. Most of the existing literature on ambiguity in RE is focused on written requirement documents, and the role of ambiguity in oral NL during elicitation interviews has not been investigated thoroughly [FSG16]. Ambiguity can cause *misunderstanding situations* during elicitation interviews, where the requirements analyst does not understand the customer’s expression or interprets it incorrectly. The latter phenomenon is known as subconscious disambiguation and is one of the major causes of requirements failure [GW89]. It is difficult to identify unless the interpretation by the analyst is not *acceptable* in his or her mental framework [FSG16]. The problem of cross-domain ambiguity can be seen as a special case of subconscious disambiguation which is caused due to different domain knowledge.

2.2 Word Embeddings

Word embedding is a collective term for language modelling techniques that map each word in the vocabulary to a dense vector representation. Contrary to one-hot representation, word embedding techniques embed each word into a low-dimensional continuous space and capture its semantic and syntactic relationships [LXT⁺15]. It is based on the distributional hypothesis proposed by Harris which states that words appearing in similar linguistic contexts share similar meanings [Har54].

One of the most popular word embedding techniques is skip-gram with negative sampling (SGNS) [MSC⁺13]. It trains a shallow two-layer neural network which, given a single input word w , predicts a set of context words $c(w)$. The context for a word w_i is the set of words surrounding it in a fixed-size window, i.e. $\{w_{i-L}, \dots, w_{i-1}, w_{i+1}, \dots, w_{i+L}\}$, where L is the context-window size. Each word w is associated with vectors $u_w \in \mathbb{R}^D$ and $v_w \in \mathbb{R}^D$, called the input and output vectors respectively. If T is the number of windows in the

given corpus, then the objective of the skip-gram model is to maximize

$$\frac{1}{T} \sum_{t=1}^T \sum_{-L \leq i \leq L; i \neq 0} \log p(w_{t+i}|w_t) \quad (1)$$

In the negative sampling method, $p(w_{t+i}|w_t)$ is defined as

$$p(w_o|w_I) = \log \sigma(u_{w_I}^T v_{w_o}) + \sum_{i=1}^k \log \sigma(-u_{w_I}^T v_{w_i}) \quad (2)$$

where $w_i \sim P(w)$ and $P(w)$ is the noise distribution.

2.3 Linear Transformation

A linear transformation can be used to learn a linear mapping from one vector space to another. Its use for combining different word embedding spaces was first explored by Mikolov *et al.* who used it for bilingual machine translation [MLS13]. They used a list of word pairs $\{x_i, y_i\}_{i=1}^n$, where y_i is the translation of x_i . Then they learned a *translation matrix* W by minimizing the following loss function

$$\sum_{i=1}^n |x_i W - y_i| \quad (3)$$

This approach can also be used for aligning monolingual word embeddings. If one assumes that the meaning of most words remains unchanged, linear regression can be used to find the best rotational alignment between two word embedding spaces. Failure to properly align a word can be then used to identify a change in meaning. This is the basis for the proposed approach towards identifying cross-domain ambiguous words. Similar approaches have been used to detect linguistic variation in the meaning of a word with time and to develop ensemble word embedding models [KARPS15, MSL17].

Significant work has been done to improve the linear transformation method. Dimension-wise mean centering has been shown to improve the performance of linear transformation methods in downstream tasks [ALA16]. Xing *et al.* noticed a hypothetical inconsistency in the distance metrics used in the optimization objectives in the work of Mikolov *et al.*: dot product for training word embeddings, Euclidean distance for learning transformation matrix, and cosine distance for similarity computations [XWLL15]. It was solved by normalizing the word embeddings and by requiring the transformation matrix to be orthogonal. The optimal orthogonal transformation matrix which maps X to Y can be found through the solution of the well-known Orthogonal Procrustes problem, which is given by

$$W = UV^T \quad (4)$$

where $X^T Y = U \Sigma V^T$ is the singular value decomposition (SVD) factorization of $X^T Y$ [Sch66].

3 Related Work

Synchronic LSC detection refers to the measurement of variation of word meanings across domains or speaker communities [SHDTSIW19]. The latter has been studied by making use of the large-scale data provided by communities on online platforms such as Reddit [TF17].

Research works on cross-domain ambiguity detection have been limited to its applicability in RE. The first approach was suggested by Ferrari *et al.* (2017) who employed Wikipedia crawling and word embeddings to estimate the variation of typical CS words (e.g., code, database, windows) in other domains [FDG17]. They used Wikipedia articles to create two corpora: a CS one and a domain-specific one, replaced the target words (top-k most frequent nouns in the CS corpus) in the latter by a uniquely identifiable modified version, and trained a single language model for both corpora. Cosine similarity was then used as a metric to estimate the variation in the meaning of the target words when they are used in the specified domain. However, this approach suffers from the following drawbacks:

- the inability to identify non-CS cross-domain ambiguous words,
- the need to construct a language model for each combination of domains, and

- the need to modify the domain-specific corpus.

Their work was extended by Mishra and Sharma who applied it on various subdomains of engineering with varying corpus size [MS19]. They identified the most suitable hyperparameters for training word embeddings on corpora of three different classes: large, medium, and small, based on the number of documents. They then used the obtained results to identify a similarity threshold for ambiguous words.

Ferrari *et al.* (2018) suggested an approach based on developing word embedding spaces for each domain, and then estimating the variation in the meaning of a word by comparing the lists of its most similar words in each domain [FEG18]. This approach addressed the above-mentioned drawbacks of the previous one. It was later extended by Ferrari and Esuli, with the major contribution being the introduction of a quantitative evaluation of the approach [FE19].

An alternative approach which does not require domain-specific word embeddings was suggested by Toews and Holland [TH19]. It estimates a word’s similarity across domains through context similarity. This approach does require trained word embeddings, but they are not required to be domain-specific, which allows it to be used on small domain corpora as well. If D_1 and D_2 are two domain corpora, then the context similarity of a word w is defined as

$$simc(w) = \frac{center(c_1) \cdot center(c_2)}{\|center(c_1)\| \cdot \|center(c_2)\|} \quad (5)$$

$$center(c) = \frac{1}{|c|} \sum_{w \in c} IDF_D(w) \cdot v_w \quad (6)$$

where $c_1 \subset D_1$ and $c_2 \subset D_2$ consist of all words from sentences containing w .

4 Motivation

The motivation behind the proposed linear transformation based approach is based on the following factors:

- The approaches suggested by Ferrari *et al.* (2018) and Toews and Holland judge a word’s meaning from its local context rather than a global one. This leads them to wrongly assign a high ambiguity score to a word having distinct, yet similar, nearest words in different domains. A particular example of this problem is the high score assigned to proper names such as *Michael*; although they are near to other proper nouns in all domains, but the exact lists vary widely. Such approaches also fail in the opposite scenario in which the meaning of the nearest words themselves change. This can happen in the case of *ambiguous clusters*. For example, a lot of topics in artificial intelligence, such as neural networks and genetic algorithms, are inspired by biology. Due to this, certain words appear together in both these domains but carry different interpretations. However, the approach proposed by this paper relies on the global context rather than the local one, which resolves the issues mentioned above.
- The proposed approach can work for more than two domains as opposed to the approaches suggested by Ferrari *et al.* (2017) and Toews and Holland [FDG17, TH19].
- The approach proposed by Ferrari *et al.* (2018) assumes the meaning of the neighbouring words to be the same across domains, whereas a linear transformation based approach works on a much weaker assumption that the meaning of most words remains the same across domains.
- Schlechtweg *et al.* evaluated various synchronic LSC detection models on *SURel*, a German dataset consisting of the meaning variations from general to domain-specific corpus determined through manual annotation [HSSiW19, SHDTSIW19]. Their study found linear transformation to perform much better than other alignment techniques such as word injection (proposed by Ferrari *et al.* (2017)) and vector initialization.

5 Approach

Given a set of domains $D = \{D_1, \dots, D_n\}$, the approach requires a word embedding space S_i corresponding to each domain D_i . The first step is to align the embedding spaces (subsection 5.1) and then determine the set of target words (subsection 5.2). The final step is to assign a cross-domain ambiguity score to each target word (subsection 5.3).

5.1 Embedding Spaces Alignment

This step determines a transformation matrix M_i for each domain-specific word embedding space S_i which maps it to a unified embedding space. It uses an algorithm devised by Muromägi *et al.* [MSL17] which iteratively finds the transformation matrices M_1, M_2, \dots, M_n and the common target space Y . It performs the following two steps in each iteration:

1. The transformation matrices M_1, M_2, \dots, M_n are calculated using equation 4.
2. The target space is updated to be the average of all transformed spaces:

$$Y(w) = \frac{1}{n_w} \sum_{i=1}^{n_w} S_i(w) M_i \quad (7)$$

where n_w is the number of domain-specific embedding spaces with word w as part of its vocabulary.

These steps are repeatedly performed as long as the change in average normalised residual error, which is given by

$$\frac{1}{n} \sum_{i=1}^n \frac{\|S_i M_i - Y\|}{\sqrt{|S_i| \cdot d}} \quad (8)$$

is equal to or greater than a predefined threshold τ .

5.2 Target Words Selection

The approach for identifying the set of target words T_D has been presented in Algorithm 1.

Algorithm 1 Algorithm for selecting target words

```

procedure SELECTWORDS( $C, k, \rho$ )
   $T_D \leftarrow \emptyset$ 
  for  $w_i \in \text{VOCAB}(C_1) \cup \dots \cup \text{VOCAB}(C_n)$  do
    if  $\text{POS}(w_i) \in \{NN, VB, ADJ\}$  then
       $\text{counts} = \{\text{FREQ}(C_1, w_i), \dots, \text{FREQ}(C_n, w_i)\}$ 
       $c_1, c_2 \leftarrow \text{TOP2VALUES}(\text{counts})$ 
      if  $c_1 \geq k \wedge c_2 \geq \rho \times c_1$  then
         $T_D \leftarrow T_D \cup \{w_i\}$ 
  return  $T_D$ 

```

This step requires two numerical parameters, k and ρ . To be considered a target word, w must satisfy three conditions:

- It must be a content word, i.e. noun, verb, or adjective.
- Its maximum frequency in a domain corpus, i.e. $f_{max} = \max(\text{count}_i(w))$, should be greater than or equal to k .
- It should have a frequency of at least ρf_{max} in any other domain corpus.

5.3 Cross-Domain Ambiguity Ranking

This step assigns an *ambiguity score* to each word in T_D based on their cross-domain ambiguity across the corpora $C = \{C_1, \dots, C_n\}$. The algorithm for the same is reported in Algorithm 2.

The idea is as follows. For each word w in the set of target words T_D , the cosine distance for each unordered pair of its transformed embeddings is calculated, which is given by

$$\text{cosineDistance}(v_i, v_j) = 1 - \frac{v_i \cdot v_j}{\|v_i\| \|v_j\|} \quad (9)$$

The average of all these cosine distances, weighted by the sum of the word frequencies in the corresponding domain corpora, is the ambiguity score assigned to the word w . All words in T_D are sorted according to their score and a ranked list A_D is produced.

Algorithm 2 Algorithm for assigning ambiguity scores

```
procedure ASSIGNAMBIGUITYSCORES( $T_D, M, S$ )  
   $Score \leftarrow \emptyset$   
  for  $w \in T_D$  do  
     $V \leftarrow \emptyset$   
    for  $S_i \in S$  do  
      if  $w \in S_i$  then  
         $V \leftarrow V \cup \{M_i S_i(w)\}$   
     $U \leftarrow 0$   
     $C \leftarrow 0$   
    for  $v_i \in V$  do  
      for  $v_j \in V \setminus v_i$  do  
         $c \leftarrow count_i(w) + count_j(w)$   
         $U \leftarrow U + c \times \text{COSINEDISTANCE}(v_i, v_j)$   
         $C \leftarrow C + c$   
     $Score[w] \leftarrow U/C$   
   $A_D \leftarrow \text{SORT}(T_D, Score)$   
  return  $A_D$ 
```

6 Results

6.1 Project Scenarios

To showcase the working of the proposed approach, this paper considers the same hypothetical project scenarios that were used by Ferrari and Esuli [FE19]. They involve five domains: computer science (CS), electronic engineering (EE), mechanical engineering (ME), medicine (MED), and sports (SPO).

- *Light Controller* [CS, EE]: an embedded software for room illumination system
- *Mechanical CAD* [CS, ME]: a software for designing and drafting mechanical components.
- *Medical Software* [CS, MED]: a disease-prediction software.
- *Athletes Network* [CS, SPO]: a social network for athletes.
- *Medical Device* [CS, EE, MED]: a fitness tracker connected to a mobile app
- *Medical Robot* [CS, EE, ME, MED]: a computer-controlled robotic arm used for surgery.
- *Sports Rehab Machine* [CS, EE, ME, MED, SPO]: a rehabilitation machine targeted towards athletes.

The first four scenarios can be thought of as an interview between a requirements analyst with a CS and a domain expert, whereas the other three scenarios can be regarded as group elicitation meetings involving stakeholders from multiple domains.

6.2 Experimental Setup

The Wikipedia API for Python¹ was used to construct the domain corpora by scraping articles belonging to particular categories. A maximum subcategory depth of 3 and a maximum article limit of 20,000 was set while creating each domain corpus.² Each article text was converted to lowercase and all non-alphanumeric words and stop words were removed, followed by lemmatization. The article count, word count, and vocabulary size for each domain corpus are reported by Table 1.

¹<https://pypi.org/project/wikipedia/>

²Since Category:Computer science is a subcategory of Category:Electronic engineering, it was excluded while creating the EE corpus to avoid extensive overlap with the CS corpus.

Table 1: Domain Corpora Statistics

Domain	Articles	Words	Vocabulary
Computer science	20,000	80,37,521	1,77,764
Electronic engineering	16,420	77,10,843	1,79,898
Mechanical engineering	20,000	1,02,02,205	1,99,696
Medicine	20,000	80,45,379	2,00,266
Sports	20,000	94,48,453	2,42,583

The word embeddings were trained using the `gensim`³ implementation of the `word2vec` SGNS algorithm with word embedding dimension $d = 50$, context window size $L = 10$, negative sampling size $\eta = 5$, and minimum frequency $f_{min} = 10$. These hyperparameters were deliberately kept equal to the values used by Ferrarri *et al.* (2018) to ensure a fair comparison between their approach and the one proposed by this paper. Training of the word embeddings was followed by length normalization and dimension-wise mean centering. For aligning the word embedding spaces, the threshold τ was set to 0.001. The plot of average normalized residual errors for each project scenario is depicted in Figure 1. The parameters for identifying target words were set as $k = 1000$ and $\rho = 0.5$. These hyperparameter values were chosen by the authors through informal experiments. The source code and the domain-specific `word2vec` models can be found in the project repository⁴.

6.3 Cross-Domain Ambiguity Rankings

The top-10 and bottom-10 ranked terms for each project scenario are reported along with their ambiguity scores in the online Appendix 1, which is available on the project repository. In order to study the cases of disagreement between the approaches proposed by this paper and Ferrari *et al.* (2018), the top-5 words with the largest absolute differences between the assigned ranks have been reported for each scenario by Table 2. The number of target words for each project scenario have also been mentioned in parenthesis.

It can be observed that most of the cases of disagreement have a higher rank, i.e. relatively lower ambiguity score assigned by the linear transformation approach proposed by this paper. Most of such cases are proper names such as *robert*, *peter*, and *daniel*. This is because of the problems associated with local-context approaches discussed in Section 4, and the low ambiguity score given to such words by the proposed approach is in line with the expected behavior of a global-context approach.

7 Conclusion and Future Work

Ambiguous requirements are a major hindrance to successful software development and it is necessary to avoid them from the elicitation phase itself. Although this problem has been studied extensively, cross-domain ambiguity has attracted research only in recent times. This paper explores the applicability of a global-context approach, which makes use of linear transformation to map various domain-specific language models into a unified embedding space, to solve this problem. From an NLP perspective, this paper is the first attempt to apply an LSC detection method on more than two domains and also the first work to logically and empirically compare the linear transformation method with the KNN-based one proposed by Ferrari *et al.* (2018).

A major challenge in applying this work in the field of requirements engineering is the suitability of Wikipedia articles as the corpora source. The future work should be primarily directed towards identifying better corpora




Figure 1: Plot of the average normalized residual errors

³<https://radimrehurek.com/gensim/>

⁴<https://github.com/vaibhav29498/Cross-Domain-Ambiguity-Detection>

Table 2: Cases of Disagreement between the Linear Transformation approach (R_1) and the one suggested by Ferrari *et al.* (2018) (R_2)

Light Controller (986)				Mechanical CAD (1016)				Medical Software (742)			
Term	R_1	R_2	$ R_1 - R_2 $	Term	R_1	R_2	$ R_1 - R_2 $	Term	R_1	R_2	$ R_1 - R_2 $
robert	972	31	941	notion	1007	18	989	peter	741	12	729
michael	933	4	929	third	31	1008	977	thomas	727	14	713
phenomenon	971	60	911	kingdom	2	975	973	third	18	730	712
peter	902	8	894	richard	964	11	953	richard	712	19	693
wide	45	939	894	green	40	987	947	mind	707	38	669

Athletes Network (569)				Medical Device (1168)				Medical Robot (1507)			
Term	R_1	R_2	$ R_1 - R_2 $	Term	R_1	R_2	$ R_1 - R_2 $	Term	R_1	R_2	$ R_1 - R_2 $
daniel	562	1	561	peter	1164	23	1141	paul	1486	8	1478
robert	569	17	552	third	56	1162	1106	coating	1501	29	1472
main	28	537	509	richard	1160	76	1084	peter	1505	38	1467
effect	522	15	507	white	74	1150	1076	third	42	1504	1462
child	59	558	499	chess	1102	39	1063	richard	1482	22	1460

Sports Rehab Machine (1624)			
Term	R_1	R_2	$ R_1 - R_2 $
daniel	1613	3	1610
love	1619	17	1602
peter	1591	16	1575
coating	1621	66	1555
told	1616	73	1543

sources which can make this area of study more applicable for the industry. Other planned future work includes a systematic quantitative evaluation of the proposed approach, extending the approach to consider multi-word phrases, and defining an ambiguity threshold.

References

- [ALA16] Mikel Artetxe, Gorka Labaka, and Eneko Agirre. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2294, Austin, Texas, November 2016. Association for Computational Linguistics.
- [AS05] K.K. Aggarwal and Yogesh Singh. *Software Engineering*. New Age International (P) Limited, 2005.
- [FDG17] Alessio Ferrari, Beatrice Donati, and Stefania Gnesi. Detecting domain-specific ambiguities: An NLP approach based on wikipedia crawling and word embeddings. In *2017 IEEE 25th International Requirements Engineering Conference Workshops (REW)*, pages 393–399, Sep. 2017.
- [FE19] Alessio Ferrari and Andrea Esuli. An NLP approach for cross-domain ambiguity detection in requirements engineering. *Automated Software Engineering*, 26(3):559–598, Sep 2019.
- [FEG18] Alessio Ferrari, Andrea Esuli, and Stefania Gnesi. Identification of cross-domain ambiguity with language models. In *2018 5th International Workshop on Artificial Intelligence for Requirements Engineering (AIRE)*, pages 31–38, Aug 2018.
- [FSG16] Alessio Ferrari, Paola Spoleтини, and Stefania Gnesi. Ambiguity and tacit knowledge in requirements elicitation interviews. *Requirements Engineering*, 21(3):333–355, Sep 2016.
- [GW89] Donald C. Gause and Gerald M. Weinberg. *Exploring Requirements: Quality Before Design*. Dorset House Publishing Co., Inc., New York, NY, USA, 1989.
- [Har54] Zellig S. Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.

- [HSSiW19] Anna HäTTY, Dominik Schlechtweg, and Sabine Schulte im Walde. SUREl: A gold standard for incorporating meaning shifts into term extraction. In *Proceedings of the Eighth Joint Conference on Lexical and Computational Semantics (*SEM 2019)*, pages 1–8, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [KARPS15] Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. Statistically significant detection of linguistic change. In *Proceedings of the 24th International Conference on World Wide Web, WWW '15*, pages 625–635, Republic and Canton of Geneva, Switzerland, 2015. International World Wide Web Conferences Steering Committee.
- [LXT⁺15] Yitan Li, Linli Xu, Fei Tian, Liang Jiang, Xiaowei Zhong, and Enhong Chen. Word embedding revisited: A new representation learning and explicit matrix factorization perspective. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, pages 3650–3656. AAAI Press, 2015.
- [MLS13] Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. *ArXiv*, abs/1309.4168, 2013.
- [MS19] S. Mishra and A. Sharma. On the use of word embeddings for identifying domain specific ambiguities in requirements. In *2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)*, pages 234–240, Sep. 2019.
- [MSC⁺13] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, pages 3111–3119, USA, 2013. Curran Associates Inc.
- [MSL17] Avo Muromägi, Kairit Sirts, and Sven Laur. Linear ensembles of word embedding models. In *Proceedings of the 21st Nordic Conference on Computational Linguistics*, pages 96–104, Gothenburg, Sweden, May 2017. Association for Computational Linguistics.
- [Pre10] Roger S. Pressman. *Software Engineering: A Practitioner's Approach*. McGraw-Hill, 2010.
- [Sch66] Peter H. Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, Mar 1966.
- [SHDTSIW19] Dominik Schlechtweg, Anna HäTTY, Marco Del Tredici, and Sabine Schulte Im Walde. A wind of change: Detecting and evaluating lexical semantic change across times and domains. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 732–746, 01 2019.
- [SM12] Yogesh Singh and Ruchika Malhotra. *Object-Oriented Software Engineering*. PHI Learning, 2012.
- [TF17] Marco Del Tredici and Raquel Fernández. Semantic variation in online communities of practice. In *IWCS 2017 - 12th International Conference on Computational Semantics - Long papers*, 2017.
- [TH19] Daniel Toews and Leif Van Holland. Determining domain-specific differences of polysemous words using context information. In *Joint Proceedings of REFSQ-2019 Workshops, Doctoral Symposium, Live Studies Track, and Poster Track co-located with the 25th International Conference on Requirements Engineering: Foundation for Software Quality (REFSQ 2019), Essen, Germany, March 18th, 2019.*, 2019.
- [WMGWF13] Yue Wang, Irene L. Manotas Gutiérrez, Kristina Winbladh, and Hui Fang. Automatic detection of ambiguous terminology for software requirements. In *Natural Language Processing and Information Systems*, pages 25–37, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [XWLL15] Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011, Denver, Colorado, May–June 2015. Association for Computational Linguistics.