Computational modeling of semantic change

Pierluigi Cassotti [1], **Francesco Periti** [2], **Stefano de Pascale** [3], **Haim Dubossarsky** [4], **and Nina Tahmasebi** [1]
[1] University of Gothenburg, [2] University of Milan, [3] KU Leuven/VUB
[4] Queen Mary University of London
{nina.tahmasebi,pierluigi.cassotti}@gu.se
francesco.periti@unimi.it, stefano.depascale@kuleuven.be
h.dubossarsky@qmul.ac.uk

1 Introduction

Languages change constantly over time, influenced by social, technological, cultural and political factors that affect how people express themselves. In particular, words can undergo the process of semantic change, which can be subtle yet significantly impact the interpretation of texts. For example, the word *terrific* used to mean "causing terror" and was as such synonymous to *terrifying*. Nowadays, speakers use the word in the sense of "excessive" and even "amazing".

In Historical Linguistics, tools and methods have been developed to analyse this phenomenon, including systematic categorisations of the types of change, the causes and the mechanisms underlying the different types of change. However, traditional linguistic methods, while informative, are often based on small, carefully curated samples. Thanks to the availability of both large diachronic corpora, the tools to model word meaning using unsupervised computational methods, and evaluation benchmarks, we are seeing an increasing interest in the computational modelling of semantic change. This is evidenced by the increasing number of publications in this new domain as well as the organisation of initiatives and events related to this topic, such as the yearly workshop on Computational Approaches to Historical Language Change *LChange*[1] that reached its fourth year , and several evaluation campaigns (Schlechtweg et al., 2020a; Basile et al., 2020b; Kutuzov et al.; Zamora-Reina et al., 2022).

Relevance Computational modelling of semantic change is highly relevant for fields like lexicography but also studies in (Historical) Linguistics where we can complement and verify existing research on larger corpora, more genres, more extended periods and many more languages. Computational modelling of semantic change is also interesting for any text-based humanities and social sciences as well as technical and medical science, where the evolution of concepts or the progression of before and after is studied. In the past few years, we have seen an increasing interest in utilizing methods for semantic change in other domains. Marjanen et al. (2019) delved into the connections between "isms" (like liberalism, socialism, and conservatism) and ideological language, shedding light on the progression of political language throughout history. Bizzoni et al. (2020) investigate changes in scientific writing, while Haider and Eger (2019) direct their focus in poetry studies. Wevers (2019) and Garg et al. (2018) investigated the presence and evolution of gender biases and ethnic stereotypes in various textual data. Vylomova et al. (2019) honed in on the semantic transformations of harm-related concepts within psychology. Their study sought to determine if concepts like *addiction*, *bullying*, *harassment*, *prejudice*, and *trauma* have broadened in scope over the past forty years. Tripodi et al. (2019) traced the evolution and prevalence of antisemitic biases across various domains, such as religion, economics, and socio-politics. Their data suggested an alarming rise in antisemitism, particularly in France, from the mid-80s onward.

This tutorial will be interest of for the ACL community as a venue for facilitating discussions and sharing knowledge on Diachronic Linguistics and time-aware language analysis. There is an extensive collection of models, methods and trained diachronic resources that benefit anyone interested in temporally evolving information beyond the LSC community. Moreover, it will provide a practical demonstration of available tools to researchers and practitioners working on different aspects of LSC and historical linguistics. In particular, we

[1] https://www.changeiskey.org/event/2023-emnlp-lchange/

Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics
Tutorial Abstracts, pages 1–8
March 21, 2024

will showcase the benchmark developed within the Change is Key! program, in which a suit of pre-trained models, as well as training and test data, are available[2], and *integrate hands-on sessions throughout the tutorial*.

2 Tutorial overview

This tutorial will overview the current approaches, problems, and challenges in detecting lexical semantic changes. At its core, the computational modelling of semantic change consists of the following:

- Modelling of word meaning, typically using unsupervised methods applied to diachronic corpora;
- modelling of meaning change, based on the outcome of the above; and
- evaluation.

This tutorial will extend the above with an introduction to lexical semantic change and an overview of the available resources (corpora, pre-trained diachronic models, and data sets). We will highlight issues in the creation and use of diachronic corpora and different procedures for annotating data. Next, we will introduce the current state-of-the-art approaches for automatic detection of LSC, provide a hands-on section on available systems and tools, and open the floor to discuss possible applications.

3 Outline

1. Introduction to Semantic Change and Computational modeling (1.5 hours)

2. Evaluation: Tasks, benchmarks, and measurements of Lexical Semantic Change (1.5 hours)

3. Models for Lexical Semantic Change Detection (2 hours)

4. Hands-on and Discussion (1 hours)

3.1 Introduction to Semantic Change and Computational modelling (1.5 hour)

We will provide a theoretical background of LSC, paying attention to semasiological phenomena, i.e., semantic change. We will introduce the classical types of semasiological change (e.g., metaphorization/metonymization or generalization/specialization) but also focus on types of changes at the level of synonymous groups or entire lexical fields (Geeraerts, 2020). Several theories, among which diachronic prototype semantics (Geeraerts, 1997) and grammaticalization theory (Traugott, 2017), will be reviewed. Finally, we will discuss some of the theoretically relevant findings recently studied in computational semantic change (e.g., the Law of Parallel Change and the Law of Differentiation (Hamilton et al., 2016a; Liétard et al., 2023; Stern, 1921)).

3.2 Evaluation: Tasks, benchmarks, and measurements of Lexical Semantic Change (1.5 hour)

We will briefly overview some of the available most used diachronic corpora such as The New York Times corpus (Sandhaus, 2008), l'Unità corpus (Basile et al., 2020a), the DTA corpus (Textarchiv), the BZ and ND corpora (Zeitung), the CCOHA corpus (Alatrash et al.), the LatinISE corpus (McGillivray and Kilgarriff, 2013), and the KubHist corpus (Adesam et al., 2019). A list of lexicographic resources useful for Lexical Semantic Change will be described, such as the Oxford English Dictionary[3] and the Sabatini Coletti dictionary[4] (Basile et al.).

We will introduce the framework DUREL (Schlechtweg et al., 2018) for the annotation of LSC, which is employed in the annotation process of Semeval 2020 Task 1 (Schlechtweg et al., 2020a). We will present the tasks on which LSC is usually framed: Unsupervised Lexical Semantic Change Detection, Lexical Semantic Change Discovery and Temporal Analogies. For each task, we will introduce the most used benchmarks, namely SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection (Schlechtweg et al., 2020b), which is the first task on Unsupervised Lexical Semantic Change Detection in English, German, Swedish, and Latin languages, RuShiftEval (Kutuzov and Pivovarova, 2021) for the Russian language, LSCDiscovery (Zamora-Reina et al., 2022), the Shared Task on Semantic Change Discovery and Detection in Spanish, NorDiaChange (Kutuzov et al., 2022), ChiWUG (Chen et al., 2023), and the datasets for the Temporal Analogies task (Yao et al., 2018; Szymanski, 2017).

[2] https://github.com/ChangeIsKey/LSCDBenchmark

[3] https://www.oed.com/

[4] https://dizionari.corriere.it/dizionario_italiano/

3.3 Models for Lexical Semantic Change Detection (2 hours)

We will provide some background on Distributional Semantics introducing PPMI matrices (Levy and Goldberg), Word2vec (Mikolov et al., 2013) and BERT models (Devlin et al., 2018). Then, we will present models for Lexical Semantic Change, starting from Alignment Models (Tahmasebi et al., 2021; Kutuzov et al., 2018; Cassotti et al., 2020). In particular, we will introduce Post-alignment models such as those based on Orthogonal Procrustes (Hamilton et al., 2016b), Jointly Explicit Alignment Models such as Dynamic word embeddings (Yao et al., 2018), and Jointly Implicit Alignment Models such as Temporal Word Embedding with a Compass (Carlo et al., 2019), Temporal Referencing (Dubossarsky et al., 2019) and Temporal Random Indexing (Basile et al., 2016).

With the increasing use of contextualised word embeddings, numerous approaches employing BERT-base models have been developed for LSC Detection (Montanelli and Periti, 2023; Laicher et al., 2021). We will present the approaches based on contextualised word embeddings following the classification framework proposed by Montanelli and Periti (2023). In particular, we will discuss the use of contextualised embeddings according to three dimensions of analysis: meaning representation, time-awareness, and learning modality. We will illustrate existing approaches as concrete examples for each dimension, allowing for a more precise and comprehensive understanding. For example, we will introduce simple unsupervised approaches such as the use of similarity measure like Average Pairwise Distance (Giulianelli et al., 2020), or clustering algorithms like WiDiD (Periti et al., 2022), but also supervised approaches that leverage the time information of the corpora such as TempoBERT (Rosin et al., 2022) and Temporal Attention (Rosin and Radinsky, 2022)).

Moreover, we will present approaches that train BERT models on Word Sense Disambiguation (Navigli, 2009) and Word-in-Context (Pilehvar and Camacho-Collados, 2019) tasks to perform LSC Detection such as GlossReader (Rachinskiy and Arefyev, 2021), DeepMistake (Arefyev et al., 2021), and XL-LEXEME (Cassotti et al., 2023). Finally, we will look at models based on lexical substitution, such as Card (2023) and Liétard et al. (2023), and generative models (Giulianelli et al., 2023).

4 Tutorial Information

Type of the tutorial Introductory.

Length This is a 6-hour tutorial.

Target audience and background This tutorial targets researchers at different levels of expertise in the field. Introductory researchers will gain a comprehensive understanding of the topic, covering foundational concepts and available resources. Intermediate researchers will deepen their knowledge with advanced approaches for automatic detection and analysis of LSC, while advanced researchers will explore state-of-the-art techniques and address complex challenges. The tutorial is designed to be inclusive, fostering the participation of attendees with varying experience levels. Furthermore, the tutorial aims to foster a more powerful synergy between the LSC domain and other areas of NLP, particularly emphasising the integration with Lexical Semantics and research pursuits in Word Sense Discrimination. Prerequisites include a basic understanding of linguistics, Natural Language Processing, and Computational Linguistics concepts.

Breadth The tutorial sections will cover both works from the tutorial presenters and others:

- Introduction to Language Change: 20% of work by tutorial presenters and 80% by others
- Evaluation: Tasks, benchmarks, and measurements of Lexical Semantic Change: 40% of work by tutorial presenters and 60% by others
- Models for Lexical Semantic Change Detection: 20% of work by tutorial presenters and 80% by others

Diversity The tutorial brings together a diverse group of presenters, each with unique computer science and linguistics backgrounds, hailing from different institutions such as the University of Gothenburg, the Queen Mary University of London, the University of Milan and Vrije Universiteit Brussel. This diverse group of experts reflects the interdisciplinary nature of the research field, where knowledge from linguistic analysis and computational methodologies converge. Furthermore, the tutorial will showcase the rich linguistic diversity of studying LSC, covering several languages, including Russian, English, Swedish, Latin, Spanish, and Italian. Exploring multiple languages will give attendees insights into how semantic change manifests across language families, historical periods, and socio-cultural contexts. The tutorial aims to

foster a global perspective on the diachronic change of word meanings by encompassing various languages, encouraging participants to draw parallels and distinctions between languages.

Audience size The proposed tutorial is expected to attract around 100+ attendees, motivated by the considerable interest and attendance observed in related events like the International Workshop on Computational Approaches to Historical Language Change and the Ever Evolving NLP (EvoNLP) Workshop.

Venue We prefer ACL 2024 and NAACL 2024 as our tutorial is tailored for an audience that includes linguists and computer scientists. EMNLP 2024 stands as our second preferred option. Should there be no available slots, we would consider EACL 2024.

Pedagogical material All materials, including presentations and Python notebooks, will be available online at the tutorial website: https://www.changeiskey.org/event/2024-eacl-tutorial/.

Past tutorials

- LREC 2022 Tutorial *Lexical Semantic Change: Models, Data and Evaluation*: While this tutorial primarily devoted its attention to resources for LSC Detection, our proposed tutorial aims to provide more comprehensive coverage on the subject of Computational Modeling of Semantic Change, as we will delve into a rich introduction of the linguistic aspects of semantic change, and a detailed exploration of computational models, emphasizing not just the conventional approaches, but also focusing extensively on the architectures of cutting-edge models.

5 Reading list

- Introduction to Semantic Change (Geeraerts et al., 2012; Traugott, 2017; Geeraerts, 2020)
- Surveys (Kutuzov et al., 2018; Tahmasebi et al., 2021; Montanelli and Periti, 2023)
- Benchmarks (Schlechtweg et al., 2020a; Basile et al., 2020c; Kutuzov and Pivovarova, 2021)
- Models (Hamilton et al., 2016c; Yao et al., 2018; Giulianelli et al., 2023; Cassotti et al., 2023; Periti et al., 2023)

6 Presenters

Nina Tahmasebi is an associate professor at the University of Gothenburg. She has researched computational methods for semantic change since 2008 and leads the *Change is Key!* program, a 6-year research program aimed at developing state-of-the-art methods for semantic change and use these to address research questions from historical linguistics as well as the humanities and social sciences. She is the chair of the LChange workshop series on Computational modeling for language change and has extensive experience in modeling and evaluation for semantic change.

Pierluigi Cassotti is a PhD student at the University of Bari (Italy) and a researcher at the University of Gothenburg (Sweden). He has been a co-organiser of the LREC 2022 Tutorial *Lexical Semantic Change: Models, Data and Evaluation*, a co-organiser of the *(LChange'23) Workshop*, and a co-organiser of the *DIACR-Ita shared task for the Italian language*. His research aims to fill the gap between Natural Language Processing tools and Diachronic Linguistics, focusing on developing models for LSCD and creating resources for the diachronic analysis of language.

Francesco Periti is a PhD student at the University of Milan (Italy). His research primarily centers around computational modeling of language change, with a specific focus on Lexical Semantic Change detection. He has been a co-organiser of the 4th International Workshop on Computational Approaches to Historical Language Change 2023 (LChange'23).

Stefano De Pascale is postdoctoral scholar at the KU Leuven (Belgium), as a member of the *Change is Key!* program, and assistant professor in Italian linguistics at the Vrije Universiteit Brussel (Belgium). He obtained his PhD in Linguistics in 2019 at the KU Leuven. In his dissertation he investigated the contribution of token-based vector space models in the study of lexical variation. In 2021 he obtained a junior FWO-postdoctoral fellowship to work on the computational modelling of diachronic prototype semantics.

Haim Dubossarsky is a lecturer for NLP at Queen Mary University of London. In his work, Haim emphasises the importance of careful methodological routines in using computational

methods in NLP as a condition for reliable and validated scientific conclusions, and is a well-cited author in the field.

References

Yvonne Adesam, Dana Dannells, and Nina Tahmasebi. 2019. Exploring the Quality of the Digital Historical Newspaper Archive KubHist. In *Proceedings of the Digital Humanities in the Nordic Countries 4th Conference, DHN 2019, Copenhagen, Denmark, March 7-9, 2019.*, CEUR Workshop Proceedings. CEUR-WS.org.

Reem Alatrash, Dominik Schlechtweg, Jonas Kuhn, and Sabine Schulte im Walde. CCOHA: Clean corpus of historical american english. In *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020*, pages 6958–6966. European Language Resources Association.

Nikolay Arefyev, Daniil Homskiy, Maksim Fedoseev, Adis Davletov, Vitaly Protasov, and Alexander Panchenko. 2021. DeepMistake: Which Senses are Hard to Distinguish for a WordinContext Model. In *Computational Linguistics and Intellectual Technologies - Papers from the Annual International Conference "Dialogue" 2021*, volume 2021-June. Section: 20.

Pierpaolo Basile, Annalina Caputo, Tommaso Caselli, Pierluigi Cassotti, and Rossella Varvara. 2020a. A Diachronic Italian Corpus based on "L'Unità". In *Proceedings of the Seventh Italian Conference on Computational Linguistics, CLiC-it 2020*, volume 2769 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Pierpaolo Basile, Annalina Caputo, Tommaso Caselli, Pierluigi Cassotti, and Rossella Varvara. 2020b. Diacr-ita @ EVALITA2020: overview of the EVALITA2020 diachronic lexical semantics (diacr-ita) task. In *Proceedings of the Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2020), Online event, December 17th, 2020*, volume 2765 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Pierpaolo Basile, Annalina Caputo, Tommaso Caselli, Pierluigi Cassotti, and Rossella Varvara. 2020c. DIACR-Ita @ EVALITA2020: Overview of the EVALITA2020 Diachronic Lexical Semantics (DIACR-Ita) Task. In *Proceedings of the Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2020)*, volume 2765 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Pierpaolo Basile, Annalina Caputo, and Giovanni Semeraro. 2016. Temporal Random Indexing: a Tool for Analysing Word Meaning Variations in News. In *Proceedings of the First International Workshop on Recent Trends in News Information Retrieval co-located with 38th European Conference on Information Retrieval (ECIR 2016)*, volume 1568 of *CEUR Workshop Proceedings*, pages 39–41, Padua, Italy. CEUR-WS.org.

Pierpaolo Basile, Giovanni Semeraro, and Annalina Caputo. Kronos-it: a dataset for the italian semantic change detection task. In *Proceedings of the Sixth Italian Conference on Computational Linguistics*, volume 2481 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Yuri Bizzoni, Stefania Degaetano-Ortlieb, Peter Fankhauser, and Elke Teich. 2020. Linguistic variation and change in 250 years of english scientific writing: A data-driven approach. *Frontiers in Artificial Intelligence*, 3.

Dallas Card. 2023. Substitution-based semantic change detection using contextual embeddings. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 590–602, Toronto, Canada. Association for Computational Linguistics.

Valerio Di Carlo, Federico Bianchi, and Matteo Palmonari. 2019. Training Temporal Word Embeddings with a Compass. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI*, pages 6326–6334, Honolulu, Hawaii,USA. AAAI Press.

Pierluigi Cassotti, Pierpaolo Basile, Marco de Gemmis, and Giovanni Semeraro. 2020. Analyzing gaussian distribution of semantic shifts in lexical semantic change models. *IJCoL. Italian Journal of Computational Linguistics*, 6(6-2):23–36.

Pierluigi Cassotti, Lucia Siciliani, Marco DeGemmis, Giovanni Semeraro, and Pierpaolo Basile. 2023. XL-LEXEME: WiC pretrained model for cross-lingual LEXical sEMantic changE. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1577–1585, Toronto, Canada. Association for Computational Linguistics.

Jing Chen, Emmanuele Chersoni, Dominik Schlechtweg, Jelena Prokic, and Chu-Ren Huang. 2023. ChiWUG: A Graph-based Evaluation Dataset for Chinese Lexical Semantic Change Detection. In *Proceedings of the 4th Workshop on Computational Approaches to Historical Language Change*, pages 93–99, Singapore. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.

Haim Dubossarsky, Simon Hengchen, Nina Tahmasebi, and Dominik Schlechtweg. 2019. Time-Out: Temporal Referencing for Robust Modeling of Lexical Semantic Change. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Volume 1: Long Papers*, pages 457–470, Florence, Italy. Association for Computational Linguistics.

Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.

Dirk Geeraerts. 1997. *Diachronic Prototype Semantics: A Contribution to Historical Lexicology*. Clarendon Press, Oxford.

Dirk Geeraerts. 2020. Semantic change. "what the smurf?". In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann, and Thomas E. Zimmermann, editors, *The Wiley Blackwell Companion to Semantics*. Wiley Blackwell, Hoboken NJ.

Dirk Geeraerts, Caroline Gevaert, and Dirk Speelman. 2012. How anger rose: Hypothesis testing in diachronic semantics. In Kathryn Allan and Justyna Robinson, editors, *Current methods in historical semantics*, pages 73–109. De Gruyter Mouton, Berlin/New York.

Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. Analysing lexical semantic change with contextualised word representations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973, Online. Association for Computational Linguistics.

Mario Giulianelli, Iris Luden, Raquel Fernandez, and Andrey Kutuzov. 2023. Interpretable word sense representations via definition generation: The case of semantic change analysis. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3130–3148, Toronto, Canada. Association for Computational Linguistics.

Thomas Haider and Steffen Eger. 2019. Semantic change and emerging tropes in a large corpus of New High German poetry. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 216–222, Florence, Italy. Association for Computational Linguistics.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016a. Diachronic word embeddings reveal statistical laws of semantic change. In *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Long Papers*, volume 3, pages 1489–1501.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016b. Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, Volume 1: Long Papers*, Berlin, Germany. The Association for Computer Linguistics.

William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016c. Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1489–1501, Berlin, Germany. Association for Computational Linguistics.

Andrey Kutuzov, Lilja Øvrelid, Terrence Szymanski, and Erik Velldal. 2018. Diachronic word embeddings and semantic shifts: a survey. In *Proceedings of the 27th International Conference on Computational Linguistics, COLING 2018*, pages 1384–1397, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Andrey Kutuzov and Lidia Pivovarova. 2021. RuShiftEval: A Shared Task on Semantic Shift Detection for Russian. In *Proc. of the International Conference on Computational Linguistics and Intellectual Technologies (Dialogue)*, 20, (online). Redkollegija sbornika.

Andrey Kutuzov, Lidia Pivovarova, and others. RuShiftEval: a shared task on semantic shift detection for russian. In *Computational linguistics and intellectual technologies: Papers from the annual conference Dialogue.* Redkollegija sbornika.

Andrey Kutuzov, Samia Touileb, Petter Mæhlum, Tita Ranveig Enstad, and Alexandra Wittemann. 2022. Nordiachange: Diachronic semantic change dataset for norwegian. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference, LREC 2022, Marseille, France, 20-25 June 2022*, pages 2563–2572. European Language Resources Association.

Severin Laicher, Sinan Kurtyigit, Dominik Schlechtweg, Jonas Kuhn, and Sabine Schulte im Walde. 2021. Explaining and improving BERT performance on lexical semantic change detection. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 192–202, Online. Association for Computational Linguistics.

Omer Levy and Yoav Goldberg. Neural Word Embedding as Implicit Matrix Factorization. In *Proc of NeurIPS*, volume 27. Curran Associates, Inc.

Bastien Liétard, Mikaela Keller, and Pascal Denis. 2023. A tale of two laws of semantic change: Predicting synonym changes with distributional semantic models. *CoRR*, abs/2305.19143.

Jani Marjanen, Lidia Pivovarova, Elaine Zosa, and Jussi Kurunmäki. 2019. Clustering ideological terms in

historical newspaper data with diachronic word embeddings. In *5th International Workshop on Computational History, HistoInformatics@TPDL 2019, Oslo, Norway, September 12, 2019*, volume 2461 of *CEUR Workshop Proceedings*, pages 21–29. CEUR-WS.org.

Barbara McGillivray and Adam Kilgarriff. 2013. Tools for historical corpus research, and a corpus of Latin. In *New Methods in Historical Corpus Linguistics*, pages 247–257, Tübingen. Narr.

Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013, Workshop Track Proceedings*.

Stefano Montanelli and Francesco Periti. 2023. A Survey on Contextualised Semantic Shift Detection. *arXiv preprint arXiv:2304.01666*.

Roberto Navigli. 2009. Word Sense Disambiguation: A Survey. *ACM Comput. Surv.*, 41(2).

Francesco Periti, Alfio Ferrara, Stefano Montanelli, and Martin Ruskov. 2022. What is Done is Done: an Incremental Approach to Semantic Shift Detection. In *Proc. of LChange*, pages 33–43, Dublin, Ireland. ACL.

Francesco Periti, Sergio Picascia, Stefano Montanelli, Alfio Ferrara, and Nina Tahmasebi. 2023. Studying Word Meaning Evolution through Incremental Semantic Shift Detection: A Case Study of Italian Parliamentary Speeches.

Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. WiC: the Word-in-Context Dataset for Evaluating Context-Sensitive Meaning Representations. In *Proc. of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.

Maxim Rachinskiy and Nikolay Arefyev. 2021. Zeroshot Crosslingual Transfer of a Gloss Language Model for Semantic Change Detection. In *Computational Linguistics and Intellectual Technologies - Papers from the Annual International Conference "Dialogue" 2021*, volume 2021-June. Section: 20.

Guy D. Rosin, Ido Guy, and Kira Radinsky. 2022. Time Masking for Temporal Language Models. In *WSDM '22: The Fifteenth ACM International Conference on Web Search and Data Mining, Virtual Event / Tempe, AZ, USA, February 21 - 25, 2022*, pages 833–841. ACM.

Guy D. Rosin and Kira Radinsky. 2022. Temporal Attention for Language Models. *CoRR*, abs/2202.02093.

Evan Sandhaus. 2008. The new york times annotated corpus. *Linguistic Data Consortium, Philadelphia*, 6(12):e26752.

Dominik Schlechtweg, Sabine Schulte im Walde, and Stefanie Eckmann. 2018. Diachronic Usage Relatedness (DURel): A Framework for the Annotation of Lexical Semantic Change. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, Volume 2 (Short Papers)*, pages 169–174, New Orleans, Louisiana, USA. Association for Computational Linguistics.

Dominik Schlechtweg, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020a. SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation, SemEval@COLING2020*, pages 1–23, Barcelona (online). International Committee for Computational Linguistics.

Dominik Schlechtweg, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020b. SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation, SemEval@COLING2020*, pages 1–23. International Committee for Computational Linguistics.

Gustaf Stern. 1921. *Swift, swiftly, and their synonyms: A contribution to semantic analysis and theory*, volume 23. Elanders boktr.

Terrence Szymanski. 2017. Temporal Word Analogies: Identifying Lexical Replacement with Diachronic Word Embeddings. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Volume 2: Short Papers*, pages 448–453, Vancouver, Canada. Association for Computational Linguistics.

Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2021. Survey of Computational Approaches to Lexical Semantic Change Detection.

Deutsches Textarchiv. Grundlage für ein referenzkorpus der neuhochdeutschen sprache. herausgegeben von der berlin-brandenburgischen akademie der wissenschaften.

Elizabeth Closs Traugott. 2017. Semantic change. In *Oxford Research Encyclopedia of Linguistics*.

Rocco Tripodi, Massimo Warglien, Simon Levis Sullam, and Deborah Paci. 2019. Tracing antisemitic language through diachronic embedding projections: France 1789-1914. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 115–125, Florence, Italy. Association for Computational Linguistics.

Ekaterina Vylomova, Sean Murphy, and Nicholas Haslam. 2019. Evaluation of semantic change of harm-related concepts in psychology. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 29–34, Florence, Italy. Association for Computational Linguistics.

Melvin Wevers. 2019. Using word embeddings to examine gender bias in Dutch newspapers, 1950-1990. In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 92–97, Florence, Italy. Association for Computational Linguistics.

Zijun Yao, Yifan Sun, Weicong Ding, Nikhil Rao, and Hui Xiong. 2018. Dynamic Word Embeddings for Evolving Semantic Discovery. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining, WSDM 2018*, pages 673–681, Marina Del Rey, CA, USA. ACM.

Frank D. Zamora-Reina, Felipe Bravo-Marquez, and Dominik Schlechtweg. 2022. LSCDiscovery: A Shared Task on Semantic Change Discovery and Detection in Spanish. In *Proc. of the Workshop on Computational Approaches to Historical Language Change (LChange)*, pages 149–164, Dublin, Ireland. Association for Computational Linguistics (ACL).

Berliner Zeitung. Diachronic newspaper corpus published by staatsbibliothek zu berlin.