Sudden Semantic Shifts in Swedish NATO Discourse

Brian Bonafilia, Bastiaan Bruinsma, Denitsa Saynova, and Moa Johansson

Chalmers University of Technology

bbonafilia@gmail.com, {sebastianus.bruinsma, saynova, moa.johansson}@chalmers.se

Abstract

In this paper, we investigate a type of semantic shift that occurs when a sudden event radically changes public opinion on a topic. Looking at Sweden's decision to apply for NATO membership in 2022, we use word embeddings to study how the associations users on Twitter have regarding NATO evolve. We identify several changes that we successfully validate against real-world events. However, the low engagement of the public with the issue often made it challenging to distinguish true signals from noise. We thus find that domain knowledge and data selection are of prime importance when using word embeddings to study semantic shifts.

1 Introduction

A well-known adage in Natural Language Processing is that one knows a word by the company it keeps (Firth, 1957). Yet, this company does not need to be stable and can change in either the long or short term. When this happens, the word undergoes a *semantic shift*. One common way to study these semantic shifts is by using temporal – or diachronic – word embeddings.

Most semantic shifts are slow and happen over many years or decades. Examples are words such as "nice", "broadcast" and "gay" which today have a different meaning than they would have had in the nineteenth century. Yet, while such shifts occur over various decennia, other shifts are more rapid. For example, the word "hero" changed its context from "veteran" and " superman" to "frontliner" and "covidwarrior" during the COVID-19 pandemic in a matter of months (Guo et al., 2022).

The speed of semantic change depends on various factors, such as whether the word has more than one meaning or how common it is in use (Hamilton et al., 2016). Also, *sudden* semantic change can occur during high-impact events, such as abrupt political, social, or cultural changes. For example, Tahmasebi et al. (2012) notes that the meaning of the word "terrorism" changed rapidly after the events of September 11, 2001. This, combined with the knowledge that a change in the meaning of a word also changes the opinions people associate with that word (Pérez and Tavits, 2023), makes understanding such sudden shifts relevant if we wish to understand people's changing opinions during real-world events.

Here, we use word embeddings to focus on an abrupt event in the case of Sweden: the country's decision to apply for NATO membership in 2022, following the Russian invasion of Ukraine. This decision was a sudden shift and a marked change in the country's stance on foreign affairs and defense.

To study this shift, we focus on the time from September 11, 2021, to September 11, 2022, the day of the 2022 Swedish general election. We chose this period as we wished to examine how the language used around NATO changed under the assumption that NATO would be a major election issue in Sweden. To measure the semantic shifts, we use the word embeddings from a Word2Vec (Mikolov et al., 2013) model to estimate the semantic context of a set of words of interest. We then track these words over time to see if and how they changed by comparing the rank sorting of the most similar words between various periods.

From here on, this paper will proceed as follows. First, we will introduce the background to the Swedish application for NATO membership, and how it can serve as a marked and sudden change. We then introduce our data and the procedure we used for pre-processing. Following this, we discuss our methods and the findings that result from them. We end with some brief conclusions and several suggestions for further research.

2 Background

For over two hundred years, Sweden followed a self-proclaimed policy of non-alignment ("alliansfrihet") (Brommesson et al., 2022). As a result,

Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics - Student Research Workshop, pages 184–193

it did not take part in most major wars, nor became part of any military alliance during the Cold War. And while it often participated in NATO exercises (Wieslander, 2022), full membership was rarely considered. Thus, Minister for Defense Peter Hultqvist could describe a Swedish membership of NATO as unthinkable as late as November 2021 (Bolin, 2023, p.307). After the invasion of Ukraine in February 2022 though, the government changed its position. This sudden change was possible due to the support of the opposition for membership and the disengagement of most citizens on the issue (Hinnfors, 2022). As a result, the government announced its plans to join NATO on April 13 and formally applied for NATO membership on May 16, 2022.

Within this timeframe, three events are of note. First, there was the Turkish opposition to Swedish membership, rooted in that country's opposition to Sweden's support for Kurdish parties and activists (Henley and Michaelson, 2022). Second, there was a "No Confidence" vote in the Swedish House of Representatives on the future of Minister for Justice Morgan Johansson. While he survived this vote thanks to the support of Kurdish-Iranian MP Amineh Kakabaveh, in return the government had to affirm an earlier agreement made in 2021 that stated that "people from those [Kurdish] organizations coming to Sweden are not terrorists" – a line of reasoning that went straight against Turkish demands (Duxbury, 2022). Third, there was the NATO Summit that took place between 28 – 30 June, where all NATO members (Turkey included) extended a formal invitation to both Finland and Sweden to join NATO.

A final point of note is that over this period, the application to NATO membership was what Berglez (2022) calls a "hidden issue". That is, both the government and opposition aimed to - and succeeded - in drawing attention away from it and were thereby followed by most of the media. An illustration of this is that the words "alliansfrihet" and "NATO" only occurred respectively 471 and 7936 times in the main Swedish media over the period of a year around the application. Moreover, the use of both words peaks around May, after which their number drops to almost zero until the elections in September.

3 Related Work

We base our decision to use global word embeddings to capture sudden semantic shifts on a well-founded body of work. Not only are they able to capture the semantic similarity and alignment between words, but they are also able to track the shifts in the meaning of political concepts. For example, Guo et al. (2022) show that the meaning of medical words changed before and after the first outbreak of Covid-19, while Rodman (2020); Rheault and Cochrane (2020) does the same for parliamentary data, and Durrheim et al. (2023) successfully use global embeddings to measure sociological concepts such as bias.

Of note is that all these papers opt to use *global* word embeddings instead of *contextual* word embeddings (e.g. ELMo (Peters et al., 2018), BERT (Devlin et al., 2019)). While *global* word embeddings associate a single embedding vector with a word, *contextual* word embeddings assign a different vector for the same word depending on the sentence in which it appears. While this has the advantage of being able to take the context of the specific occurrence of a word into account, it does not provide a way to represent the position of a single word in the embedding space. That is, when we care about the global shift of words (as we do here), we need a global and not a contextual embedding. As such, most authors in the social sciences, and we here as well, opt to use global embeddings.

4 Data

To measure our semantic shifts, we rely on Swedish-language Twitter posts ("tweets") that focus on NATO. We do so as Twitter's broad user base touches all segments of society, allowing us to get a complete picture of the debate around NATO. Besides, as tweets have a limit of 280 characters, their length is very similar. This has the advantage that it improves data consistency while reducing computational complexity.

Within our year-long period, we collected $1,188,556$ tweets, made by a total of $64,315$ users participating in $507,359$ conversations. Of these, $329,336$ are retweets, leaving $859,220$ original tweets. We collected a tweet if it contained any one of a set of search terms relating to NATO. To generate these terms, we drew on both theoretical expectations (deductive) as well as first results (inductive). As such, we ended up with seventy-five unique search terms covering NATO, alliances, and the war in Ukraine (see Bonafilia (2023) for a complete list). Many of these words were either compound words that contain “nato” or relate to NATO and are specific enough to only occur in that context. Thus, we did not include general terms such as “allians” (alliance), unless they were part of the phrase “militär allians” (military alliance) or “allians med turkiet” (alliance with Turkey). In the end, we included a tweet when: a) it contained any of the search terms, b) the tweet is a response to another tweet that contained a search term, or c) the tweet has a response containing a search term.

Based on the background of the NATO issue as sketched above, we divide our tweets into four periods. First, there is the pre-invasion period, ranging from September 11, 2021, to 24 February 2022 (the date of the Russian military invasion of Ukraine). Second, there is the post-invasion period running from February 24 to April 13, the date of the joint press conference of the Swedish PM Andersson and her Finnish colleague Marin, where both announced the possibility of their countries joining NATO. Third, there is the pre-application period, running between April 13 and the formal application on May 16. Finally, there is the post-application period, running between May 16 and the elections on September 11. Table 1 shows the number of tweets for each of the periods.

	Tweets	Words
Pre-Invasion	131 889	2.3 M
Post-Invasion	413 517	6.8 M
Pre-Application	294 453	5.1 M
Post-Application	346 948	5.4 M

Table 1: Sizes of the Twitter dataset for each period.

To support our choice for these four periods, we look at the daily number of tweets we gathered (see Figure 1). Here, we see that at the boundaries of the four periods (indicated by arrows 2, 3, and 5) there are clear peaks in the number of tweets. Besides, we find smaller peaks between January 15 – 19 (during the Russian military build-up near the Ukrainian border), on May 13 (the first Turkish signal of opposition to Sweden’s entry into NATO), on June 7 (during the “No Confidence” vote against Morgan Johansson), and on June 28 (the NATO summit in Madrid).

5 Pre-Processing

Given that the choice – and order of – pre-processing steps will influence our analysis, we discuss each of these steps in turn (Denny and Spirling, 2018). First, we remove any URLs and mentions to other users as well as some minor punctuation. Second, we split our tweets into individual tokens. For this, we use the NLTK library’s *nltk.TweetTokenizer*, as it splits hashtags and emojis better than other tokenizers (Bird et al., 2009). Third, we lowercase all tokens, create n-grams (with no limit, so 3-grams can occur), and remove all remaining punctuation. Finally, we normalize the spelling of our tokens to address the various spellings of the same word (e.g. “grey” and “gray”). For a more detailed overview of the pre-processing see Bonafilia (2023).

We did not perform the common steps of removing stop words or lemmatizing the tokens, as we found that these steps weakened the relationship between related words. Singletons and low-frequency words were filtered out by the Gensim library (Řehůřek and Sojka, 2010), which was used for the analysis.

6 Method

The model we chose to find our word embeddings is *Word2Vec* (Mikolov et al., 2013). This is a single-layer neural network that is trained to predict a word from its context – Continuous Bag-of-Words (CBOW) – or context from a given word – Skip-gram (SG). We opted to use both architectures given that they are different in the associations they capture, their computational efficiency, and their sensitivity to less-frequent words (Mikolov et al., 2013).

6.1 Training of the Model

As with all other embedding models, *Word2Vec* needs a large amount of text to be able to capture word associations. As the tweets from each period contained insufficient data to train a new model, we used Twitter data for each period to *fine-tune* an already trained model representing general Swedish. This initial model was trained on Swedish media text (Göteborgs-Posten, SVT, and Wikipedia) from 2003 until 2014, made available by Språkbanken’s Korp language resource (Borin et al., 2012). The total number of tokens in this corpus is 759 million, with about 1.04 million unique tokens which appear at least ten times. We chose the cut-off dates

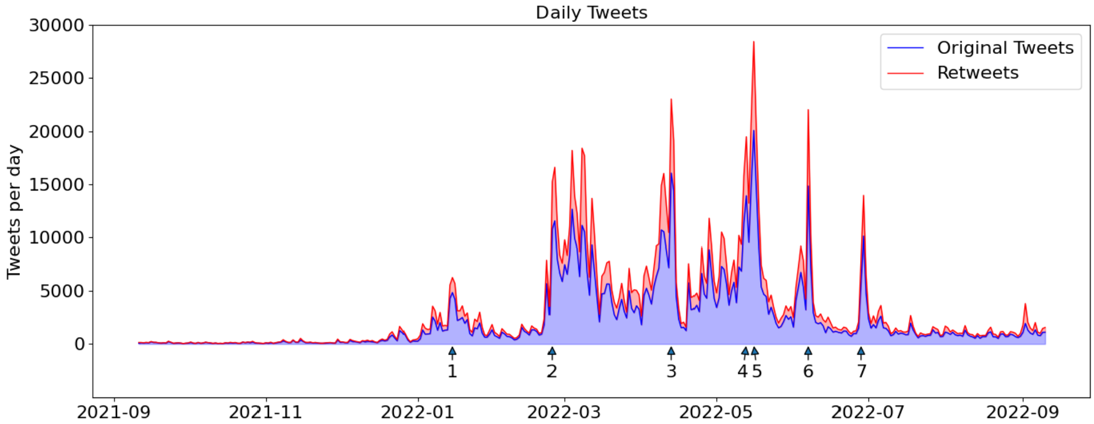

Figure 1: Tweets found by the search criteria from 2021-09-11 until 2022-09-11. The timeline of tweets with key dates: (1) January 2022, build-up of Russian forces near the Ukrainian border, (2) February 24th, the Russian Federation invades Ukraine, (3) April 13th, Swedish and Finnish PMs hold a joint press conference about the decision to join NATO, (4) Turkey expresses opposition to Sweden joining NATO on May 13th, and (5) May 16th, Sweden and Finland formally apply to join NATO. (6) A vote of No Confidence is held for Justice Minister Morgan Johansson on June 7th, and (7) the NATO Summit in Madrid on June 28th.

of 2003 and 2014 to avoid biasing the model with inputs from after the Maidan uprisings in Ukraine in 2014. This control over the input period and model parameters was our main motivation to train and validate a new model rather than use a publicly available set of pre-trained vectors.

We then trained two base models – one for the Skip-gram and one for the Continuous Bag-of-Words architecture. For both, we used Negative Sampling, a window size of 5, a minimum number of word occurrences of 10, and 160 training iterations. To validate our base model, we used the word similarities and relatedness from SuperSim by Hengchen and Tahmasebi (2021) and a QVEC-CCA scoring as introduced by Tsvetkov et al. (2016) using a Swedish pack available from Språkbanken's Korp (Borin et al., 2013). In all cases, the results indicated that the base models were well trained (Bonafilia, 2023).

We then fine-tuned both the SG and CBOW architectures on the tweets made within each period, using our pre-trained models as a base. Because the Word2Vec model training is a stochastic process, and as we have to account for instability due to data variability, we trained 10 models for each case on a different uniform random sample of 90% of the text data from that period when we perform our bootstrapping. We then ranked the most similar words based on the average cosine similarity across all 10 models.

6.2 Analysis Approach

Once we have our model, we have to formalize a search method to decide which words we want to select to look at. While we are aware that we could use the embeddings themselves to find the most similar and most different words – we opt here for a *subjective* approach. The reason for this is that we know our topic of interest (NATO) and can draw on prior knowledge not included in the model.

For the core selection of words, we take those that have either a direct relation to NATO or are synonymous with it (e.g. "försvarsalliansen" (defense alliance)), have a link to states or persons involved in Sweden's application (e.g. "erdogan", "putin", "finland"), have an association with the topics raised in the NATO discussion (e.g. "suveränitet" (sovereignty)), or words for which one subset of users in the polarization study had a markedly different use as indicated by word embeddings than another subset of users (e.g. "inkompetent" (incompetent), "dotters" (daughter's)). Besides this, we also draw on a study of words linked to polarized opinions on the issue of Sweden's entry into NATO (Bonafilia, 2023). In the end, this results in a list of 8000 words.

We then use these 8000 words and compare the averaged most similar words across the different time steps to find novel associations. While doing so, we ignore words that appeared in similar placements in all periods, such as synonyms or inflections of the word of interest. As not all the 8000 show interesting behavior, we then perform a second selection of words.

For refining the selection of words, we take all those words that fall under any one of the following criteria:

– Words which domain knowledge suggested are relevant.

– Words seen to be polarizing by Bonafilia (2023).

– Words which markedly changed their most similar words from the pre-trained model or between periods as determined by Rank-Biased Overlap (RBO) (Webber et al., 2010) of the sorted list of most similar words.

– Words for which unique words appeared among the most similar words in one of the periods but not among the most similar words in any other period.

After this second selection, we perform a last, manual review to look at general trends and to drop noisy findings. We did so as we wanted to drop those words which had very different embedding only because they were too infrequent to have a meaningful embedding at all.

7 Results

As both architectures lead to different results, we will look at both the results of the Continuous Bag-of-Words (CBOW) and the Skip-gram (SG) in turn. For each of the two, select four words that we deemed showed interesting patterns. These are "natoansökan" (NATO application) and "försvaret" (defense), as well as two unique ones for each – "nato" (NATO) and "säkerhet" (security) for COBW and "förskolor" (preschools) and "putin" (Putin) for SG. For each word, we give the top four words associated with it based on their cosine similarity. Besides these, we will also reflect on several other words that we found showed interesting behavior.

7.1 Continuous Bag-of-Words

Table 2 shows the words with the highest cosine similarity for each of the four words for the CBOW model. Also, in Figure 2, we show, for each of these four words, the comparison of the Rank-Biased Overlap between the list of the most similar words for each period and the list from the pre-trained CBOW model. Words such as "natoansökan", "nato" and "säkerhet" have a consistently low agreement in all periods, indicating a substantial shift from the base model. While "försvaret" drops to zero in the Pre-Application period as the agreement is lost completely, however, from Table 2 it is hard to determine the meaning of the shift, illustrating the difficulty in isolating the signal from noise and interpreting the results. In the pre-training data, "natoansökan" (NATO application) is so infrequently used that the word embeddings are meaningless. In the period leading up to the application, the subject of Sweden's NATO application becomes topical enough that a hashtag (#natoansökan) starts to be used. Also, for the topic of "säkerhet" (security), we find that it becomes related to the concepts of "suveränitet" (sovereignty) as the discussion of Sweden giving up neutrality to join a defensive alliance takes shape.

The word "nato" itself, becomes closely associated with the word "sverige" (Sweden), as both have a higher frequency (11×10^{-3}) and (6×10^{-3}) when compared with the pre-trained data (1×10^{-5} and 8×10^{-4} respectively). Leading to the word "nato" having a more meaningful word embedding in the base model. The reason for this is that "nato", being one of the search words, is so frequent in our data, that it has a high association with all other words. This makes the embedding relatively uninteresting to look at, as the embedding of the word is more related to other words of high frequency - such as "sverige" (Sweden) and "vi" (we) - than with words of similar meaning. This underscores the limitation of using word embeddings to find meaningful shifts for words that are deliberately sought out to generate the dataset.

7.2 Skip-gram

Table 3 shows the words with the highest cosine similarity for the Skip-gram architecture and Figure 3 shows the RBO results. Here, it can be seen that during the period after the Russian invasion of Ukraine and before the application, there is an association between "natoansökan" (NATO application) and "destabiliserande" (destabilizing). References to destabilization appeared almost exclusively during this period. This also fits well with the political consensus at the time, i.e. that a Swedish application to NATO would destabilize the country by jeopardizing its relationship with Russia. After the press conference on April 13, this changed and an association with "eventuell" (possible) and other words relating to the (likeliness of the) process of

	natoansökan	försvaret	nato	säkerhet
Base	sverigesregering	försvarsmakten	försvarsalliansen	rättssäkerhet
	regeringsbildandet	flygvapnet	fn	trovärdighet
Pre-Invasion	osansökan	försvarsmakten	sverige	säkerhetspolitik
	emuomröstning	underhållet	ukraina	natoansökning
	natooption	rättsväsendet	usa	konkurrenskraft
	intresseanmälan	välfärdssystemet	vi	stabilitet
Post-Invasion	medlemskapsansökan	försvarsmakten	sverige	säkerhetspolitik
	natoanslutning	totalförsvaret	vi	natoansökning
	dispensansökan	försvarsanslaget	ukraina	suveränitet
	osansökan	försvarsförmågan	finland	frihet
Pre-Application	#natoansökan	underhållet	sverige	suveränitet
	natomedlamskap	försvarsförmågan	#nato	rättssäkerhet
	ansökningsprocess	bnp	finland	försvarskapacitet
	medlemskapsansökan	insatsförsvaret	vi	säkerhetspolitik
Post-Application	natoanslutningen	luftsförsvaret	sverige	säkerhetspolitik
	natoprocess(en)	totalförsvaret	finland	överlevnad
	natomedlemskap	välfärdssystemet	turkiet	oljeförsörjning
	natoansökningen	insatsförsvaret	#nato	suveränitet

Table 2: Words with top cosine similarity in Continuous Bag-of-Words models grouped by period, for “natoansökan” (NATO application), “försvaret” (defense), “nato” (NATO), and “säkerhet” (security)

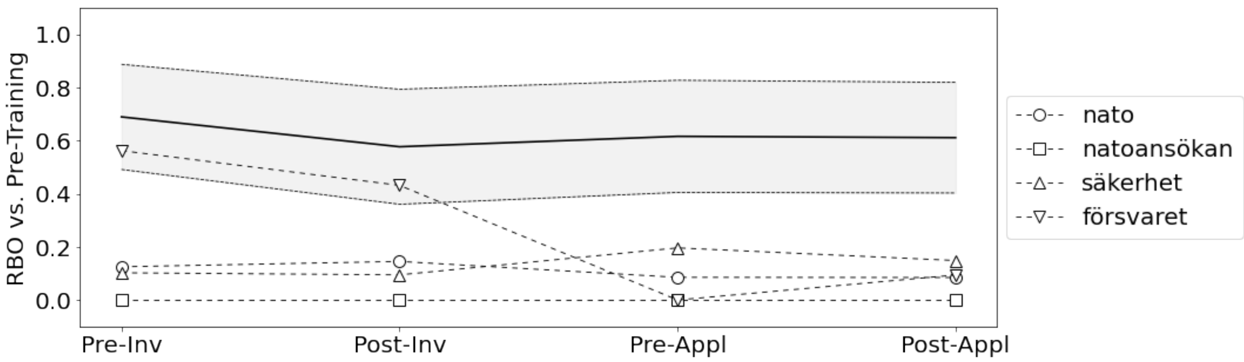

Figure 2: Comparison of the Rank-Biased Overlap between a list of the most similar words in each period and the pre-trained CBOW model for a small selection of words. A higher RBO value signifies more agreement with the base model and therefore a smaller semantic shift. The Average (solid line) and one standard deviation (shading) for 1000 randomly chosen words are also shown.

application, began to appear. We can see a similar change for “försvaret” (defense) from where the association shifts from words relating to maintenance and juridical matters before the application to a connection to the spending goal of 2% of GDP (the words “2%” and “bnp”) for NATO members afterward.

Furthermore, we see a neutral word such as “förskolor” (preschool) has a strong cosine similarity to “kärnvapen” (nuclear weapons) in the period leading up to the application. While seemingly contradictory, the reason behind this is that during this time, Left Party leader Nooshi Dadgostar made a public statement regarding not wanting NATO’s nuclear weapons to be housed within Sweden, alluding to a possibility of nuclear weapon silos near her daughter’s preschool. This generated conversation among Twitter users discussing the pros and cons of the NATO application, resulting in the SG model finding the similarity in the contexts in which these words appeared in. Also, we see the emergence of novel words related to Vladimir Putin. For ex-

	natoansökan	försvaret	förskolor	putin
Base	sverigesregering	försvarsmakten	skolor	vladimirputin
	ratificera	flygvapnet	äldreboenden	medvedev
Pre-Invasion	medlemsansökan	försvarsmakten	gymnasieskolor	ryssland
	byggförhandlingarna	invasionförsvaret	äldreboenden	biden
	omvärldsutveckling	förbandsverksamhet	fritidshem	nato
	drömregering	fm	vårdcentraler	xi
Post-Invasion	medlemsansökan	försvarsmakten	polisstationer	ryssland
	natomedlemskap	bnp	äldreboenden	han
	destabilisera(nde)	rusta	gymnasieskolor	ukraina
	natoanslutning	anslagen	fritidshem	nato
Pre-Application	natoanslutning	bnp	dotters	ryssland
	eventuell	2%	dagis	putler
	natomedlemskap	rusta	kärnvapen	erdogan
	svensk	försvarskostnaderna	kärnvapenbaser	ryssen
Post-Application	sveriges	bnp	skolbibliotek	erdogan
	finlands	2%	förskoleverksamhet	ryssland
	natoprocessen	försvarsanslaget	fritidshem	biden
	inlämnad	materielanskaffning	gymnasieskolor	putler

Table 3: Words with top cosine similarity in Skip-gram models grouped by period, for “natoansökan” (NATO application), “försvaret” (defense), “förskolor” (preschools) and “putin” (Putin)

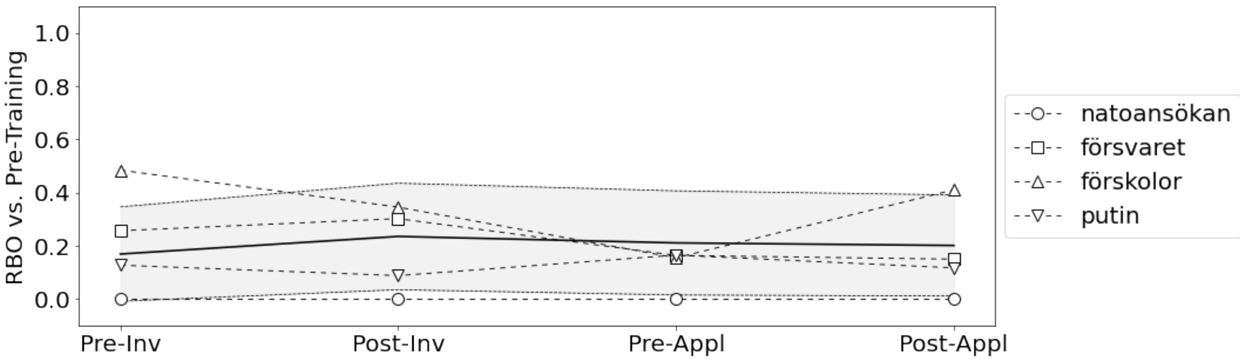

Figure 3: Comparison of the Rank-Biased Overlap between a list of the most similar words in each period and the pre-trained SG model for a small selection of words. A higher RBO value signifies more agreement with the base model and therefore a smaller semantic shift. The Average (solid line) and one standard deviation (shading) for 1000 randomly chosen words are also shown.

ample, the word “putler” is meant to draw a connection between Russia’s invasion of Ukraine and the aggression of Nazi Germany during the Second World War. Finally, when looking at the RBO results, in contrast to CBOW, SG shows a larger average shift from the baseline model for all periods. This results in the approach yielding less clear results and the need for more noise words to be filtered to find useful examples, making it harder to detect a true signal. For example, even when “förskolor” becomes a relevant word, the dip in the rank order similarity is small since the similarity was low across the board.

7.3 Further Examples

Other words (not shown here), also exhibit a strong relationship with certain events during the period. Thus, the word “inkompetent” (incompetent) first had associations with words like “korrumperad” (corrupted) and “felprioriteringar” (misplaced priorities), but later switched those to words such as “minister” (Minister), and “morganjohansson”

(Morgan Johansson) at the time of the vote of no-confidence against Minister for Justice Morgan Johansson. Besides, the word "natomotståndare" (NATO opponent), while first being associated with the Left Party (a traditional opponent of Swedish NATO membership), became associated with the Green Party and individual Social Democrats (such as former Minister for Defense Peter Hultqvist) instead. Finally, as expected, we observe that the word "kiev" is first associated with other cities, such as Tbilisi, while Post-Invasion it gains an association with the Ukrainian "kyiv" spelling, presumably by Twitter users who wished to express solidarity with Ukraine. Finally, while the word "azov" in the pre-training data referred to the Sea of Azov or any of a number of Ukrainian and Russian locations, the most similar words were other places in the area. Later, during the Post-Invasion period, this changed. First, the use of "azov" centered around the alleged neo-Nazi ties of the Azov Battalion, a Ukrainian militia, and then later became associated with the Siege of Mariupol, where defenders had occupied the "Azovstal" Steel Plant.

8 Conclusion

Our aim with this study was to look at the sudden semantic shift that we expected to occur when Sweden decided to apply for NATO membership in 2022. Looking at various words related to this application process, we find that word embeddings are a powerful tool to capture some of those shifts. Moreover, when validating them against real-world events, we find that those shifts are both accurate and meaningful. Yet, the sparsity of the dataset often makes it difficult to separate signal from noise when looking at the model results alone.

The misalignment between the signals that each of the two model architectures – SG and CBOW – manage to capture, as well as the difficulty of validating and interpreting the results exemplifies the challenges in using word embeddings for automatically detecting and measuring semantic shifts. Thus, there is a need for extensive human interpretation and validation based on domain knowledge together with a broad range of statistics that can reveal different aspects of the patterns captured by the models. Despite this though, word embeddings are still a powerful method that can aid the discovery process. As we showed, they are efficient enough to process large amounts of data and capture several underlying word relationships and sudden semantic shifts.

9 Suggestions for Further Research

We see two suggestions for further research, two methodological and one practical. On the methodological side, we saw that selecting Tweets by their relationship to NATO resulted in a skewed frequency of NATO-related words when compared with those in the pre-trained model. Such a sparse dataset with non-representative word distributions makes the study of the search words hard. To allay this, one could extend the criteria to capture a broader and more diverse representation of the language used during the period.

Another methodological option is the consideration of a different model. Two alternatives to the model we used here are FastText (Joulin et al., 2017) and GloVe (Pennington et al., 2014). Both offer a different perspective on word embeddings and might address some of the issues we faced here.

From the practical side, we assumed that Swedish NATO membership would be a major electoral issue and that a single year was enough to capture this debate. Both proved to be wrong. NATO membership was rarely discussed in the period leading up to the elections, and at the time of writing, Sweden's NATO aspirations are still unfulfilled. Thus, further research could extend the data collection period to gain a better view of any shifts in the word embeddings.

Acknowledgments

This work was supported by the Wallenberg AI, Autonomous Systems and Software Program – Humanities and Society (WASP-HS) funded by the Marianne and Marcus Wallenberg Foundation and the Marcus and Amalia Wallenberg Foundation.

References

Peter Berglez. 2022. Hur påverkades valrörelsen 2022 av omvärlden? In Niklas Bolin, Kajsa Falasca, Marie Grusell, and Lars Nord, editors, *Snabbtänkt 2.0 22*, page 111. Mittuniversitetet, Demicom, Sundsvall.

Steven Bird, Edward Loper, and Ewan Klein. 2009. *Natural Language Processing with Python*. O'Reilly Media, Sebastopol, CA.

Niklas Bolin. 2023. The Repercussions of the Russian Invasion of Ukraine on the Populist Radical Right in Sweden. In Gilles Ivaldi and Emilia Zankina, editors, *The Impacts of the Russian Invasion of Ukraine* *on Right-Wing Populism in Europe*, pages 302–313. European Center for Populism Studies (ECPS), Brussels.

Brian Bonafilia. 2023. Methods for Detecting Echo Chambers in Social Media Networks. Master's thesis, Chalmers University of Technology.

Lars Borin, Markus Forsberg, and Lennart Lönngren. 2013. SALDO: A Touch of Yin to WordNet's Yang. *Language Resources and Evaluation*, 47(4):1191–1211.

Lars Borin, Markus Forsberg, and Johan Roxendal. 2012. Korp – the corpus infrastructure of Språkbanken. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Paris. European Language Resources Association (ELRA).

Douglas Brommesson, Ann-Marie Ekengren, and Anna Michalski. 2022. Sweden's Policy of Neutrality: Success Through Flexibility? In Caroline de la Porte, Guðný Björk Eydal, Jaakko Kauko, Daniel Nohrstedt, Paul 't Hart, and Bent Sofus Tranøy, editors, *Successful Public Policy in the Nordic Countries: Cases, Lessons, Challenges*, pages 284–305. Oxford University Press, Oxford.

Matthew J. Denny and Arthur Spirling. 2018. Text Preprocessing For Unsupervised Learning: Why It Matters, When It Misleads, And What To Do About It. *Political Analysis*, 26(2):168–189.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings 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 4171–4186, Minneapolis, MN. Association for Computational Linguistics.

Kevin Durrheim, Maria Schuld, Martin Mafunda, and Sindisiwe Mazibuko. 2023. Using Word Embeddings to Investigate Cultural Biases. *British Journal of Social Psychology*, 62(1):617–629.

Charles Duxbury. 2022. Swedish Government Narrowly Survives No-Confidence Vote. Politico (EU), 06-07-2022.

J. R. Firth. 1957. Applications of General Linguistics. *Transactions of the Philological Society*, 56(1):1–14.

Yanzhu Guo, Christos Xypolopoulos, and Michalis Vazirgiannis. 2022. How COVID-19 is Changing Our Language: Detecting Semantic Shift in Twitter Word Embeddings. In *Conférence Nationale en Intelligence Artificielle 2022 (CNIA 2022)*, Saint-Etienne, France.

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

Simon Hengchen and Nina Tahmasebi. 2021. SuperSim: A Test Set for Word Similarity and Relatedness in Swedish. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 268–275, Linköping. Linköping University Electronic Press.

Jon Henley and Ruth Michaelson. 2022. Erdoğan: Turkey 'not positive' about Sweden and Finland joining Nato. The Guardian, 13-05-2022.

Jonas Hinnfors. 2022. Socialdemokraterna: högervridning och hot utifrån. In Niklas Bolin, Kajsa Falasca, Marie Grusell, and Lars Nord, editors, *Snabbtänkt 2.0 22*, page 39. Mittuniversitetet, Demicom, Sundsvall.

Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 427–431. Association for Computational Linguistics.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Stroudsburg, PA. Association for Computational Linguistics.

Efrén Pérez and Margit Tavits. 2023. *Voicing Politics: How Language Shapes Public Opinion*. Princeton University Press, Princeton, NJ.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, Stroudsburg, PA. Association for Computational Linguistics.

Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta. ELRA.

Ludovic Rheault and Christopher Cochrane. 2020. Word Embeddings for the Analysis of Ideological Placement in Parliamentary Corpora. *Political Analysis*, 28(1):112–133.

Emma Rodman. 2020. A Timely Intervention: Tracking the Changing Meanings of Political Concepts with Word Vectors. *Political Analysis*, 28(1):87–111.

Nina Tahmasebi, Gerhard Gossen, Nattiya Kanhabua, Helge Holzmann, and Thomas Risse. 2012. NEER: An Unsupervised Method for Named Entity Evolution Recognition. In *Proceedings of COLING 2012: Technical Papers*, pages 2553–2568, Mumbai. The COLING 2012 Organizing Committee.

Yulia Tsvetkov, Manaal Faruqui, and Chris Dyer. 2016. Correlation-based Intrinsic Evaluation of Word Vector Representations. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 111–115, Stroudsberg, PA. Association for Computational Linguistics.

William Webber, Alistair Moffat, and Justin Zobel. 2010. A Similarity Measure for Indefinite Rankings. *ACM Transactions on Information Systems*, 28(4):1–38.

Anna Wieslander. 2022. "The Hultqvist doctrine" – Swedish Security and Defence Policy after the Russian Annexation of Crimea. *Defence Studies*, 22(1):35–59.