

# ReFRAME or Remain: Unsupervised Lexical Semantic Change Detection with Frame Semantics

Bach Phan-Tat<sup>1</sup>, Kris Heylen<sup>1,2</sup>, Dirk Geeraerts<sup>1</sup>, Stefano De Pascale<sup>1,3</sup>, Dirk Spielman<sup>1</sup>

<sup>1</sup>Department of Linguistics, KU Leuven

<sup>2</sup>Instituut voor de Nederlandse Taal

<sup>3</sup>Vrije Universiteit Brussel

Correspondence: [ttbach.phan@kuleuven.be](mailto:ttbach.phan@kuleuven.be)

## Abstract

The majority of contemporary computational methods for lexical semantic change (LSC) detection are based on neural embedding distributional representations. Although these models perform well on LSC benchmarks, their results are often difficult to interpret. We explore an alternative approach that relies solely on frame semantics. We show that this method is effective for detecting semantic change and can even outperform many distributional semantic models. Finally, we present a detailed quantitative and qualitative analysis of its predictions, demonstrating that they are both plausible and highly interpretable.

## 1 Introduction

Lexical semantic change detection is well-established in Natural Language Processing (NLP), with different shared tasks for different languages, amongst which the most popular one is SemEval 2020 Task 1 (Schlechtweg et al., 2020). Most modern systems are based on either static word embeddings like word2vec (Mikolov et al., 2013) or ELMo (Peters et al., 2018), or contextualised ones such as BERT (Devlin et al., 2019) or S-BERT (Reimers and Gurevych, 2019), the latter currently achieving the state-of-the-art (SOTA) results (Cassotti et al., 2023). Despite their architectural differences, these methods are grounded in the distributional hypothesis (Firth, 1957), and can capture some aspects of lexical meaning (e.g., Iacobacci et al., 2016; Lenci et al., 2022; Pilehvar and Camacho-Collados, 2019; Yenicelik et al., 2020). Their application in the domain of diachronic semantics has therefore been a logical next step.

However, co-occurrence is only one of the many perspectives of lexical semantics. Especially in the cognitive-linguistic literature, many different frameworks have been proposed, such as Behavioural Profiles (Gries and Divjak, 2009), Frame Semantics (Fillmore, 1982), Conceptual Metaphor




Figure 1: Changes in the frequencies of 6 frames that the English noun *plane* participated in between the 2 subcorpora for English of the SemEval 2020 shared task 1 (Schlechtweg et al., 2020). *Plane* is annotated as semantically changed in the dataset

(Lakoff and Johnson, 1981) (see Geeraerts (2010) for a comprehensive overview of lexical semantics), and some of them have also been used in diachronic lexical semantics and semantic change research (e.g., Geeraerts, 1983; Jansegers and Gries, 2020; Law, 2019; Pettersson-Traba, 2016; Sovran, 2004; Sullivan, 2007; Sweetser, 1990; Vais, 2022).

Exploiting Frame Semantics for detecting lexical semantic change is the focus of this paper. Consider the English noun *plane*: in the 20<sup>th</sup> century, after the birth of the modern airplane in 1903, the AIRCRAFT meaning emerged and became more dominant while the DIMENSIONAL meaning became less dominant. This is reflected in the shift in the frame distributions of *plane*, where AIRCRAFT-related frames such as *commerce\_sell*, *departing*, *self\_motion* increase in frequency and DIMENSIONAL-related frames such as *identicality*, *part\_whole* decrease in frequency, as shown in Figure 1.

Our main hypothesis is that **semantic change is reflected in changes in the frames in which a word participates, either as a frame-evoking element or as a frame element**. This view is com-

patible with distributional approaches: large-scale distributional models may also capture aspects of frame structure implicitly. The contrast lies instead in the level of representation. Rather than relying on implicit information in the embedding space, we use explicitly annotated frame-semantic information as a simple and transparent signal for lexical semantic change. Our aim is therefore not to establish a new state of the art, but to draw attention to interpretable methods that make use of explicit linguistic information.

We investigate this question experimentally<sup>1</sup> using the English dataset<sup>2</sup> from SemEval 2020 Task 1. Our main findings are the following:

- Tracing the changes in the distribution of frames outperforms many distributional semantic models, including both static embeddings (count-based and prediction-based) and contextualised embeddings.
- Because the method is interpretable by design, we can decompose change scores into contributions of individual frames and inspect supporting contexts, enabling targeted qualitative analyses to understand how a word has changed (as in the plane example above) in contrast to the opaqueness of dense embedding models<sup>3</sup>.

## 2 Related work

A long-standing tradition in computational diachronic semantics builds on the Distributional Hypothesis (Firth, 1957). With the rise of corpus linguistics and concordancing, researchers began tracing a target word’s collocates and recurring patterns across corpora and time slices, often using association measures to identify characteristic collocations (Church and Hanks, 1990; Sinclair,

1991). This collocational approach has been applied in many case studies (e.g., Alba-Salas, 2007; Pettersson-Traba, 2021), methods (Tat et al., 2026; Garcia and García Salido, 2019; McEnery et al., 2019), and tools (Anthony, 2014; Brezina et al., 2020; Kilgarriff et al., 2014; Jurish, 2015). With advances in computational modelling, distributional information was represented as numerical vectors, from count-based matrices (Bullinaria and Levy, 2007; Gulordava and Baroni, 2011; Schütze, 1998) to dense neural embeddings (Baroni et al., 2014; Bojanowski et al., 2017; Levy and Goldberg, 2014; Mikolov et al., 2013; Pennington et al., 2014) and contextualised token-level models (Devlin et al., 2019). In lexical semantic change research, standard workflows typically split corpora into time slices, learn and align representations, and measure change through distances or clustering (Tahmasebi et al., 2021; Tahmasebi and Dubossarsky, 2023). As shared tasks and standardised datasets became central, compact vector representations and single change scores became especially attractive, and now dominate much of the field, including current SOTA approaches (e.g., Cassotti et al., 2023; Geeraerts et al., 2023; Giulianelli et al., 2020; Kim et al., 2014; Kulkarni et al., 2015; Kutuzov et al., 2018; Montariol et al., 2021; Tahmasebi and Dubossarsky, 2023; Tahmasebi et al., 2021).

Other theory-driven frameworks have also been employed in Diachronic Semantics (e.g., Geeraerts, 1983; Jansegers and Gries, 2020; Law, 2019; Pettersson-Traba, 2016; Sovran, 2004; Sullivan, 2007; Sweetser, 1990; Vais, 2022). However, these approaches mostly lend themselves to the description of specific linguistic phenomena, making it difficult to compare systematically across languages and datasets. Recently, efforts have been made to benchmark these theory-driven approaches with standard shared tasks (e.g., Kutuzov et al., 2021; Ryzhova et al., 2021 on grammatical profiles; Tang et al., 2023 on sense distributions). Their strong performances on shared-task benchmarks (sometimes even outperform many embedding-based methods) inspired us to test the effectiveness of using Frame Semantics information in capturing word meaning change.

Frame Semantics was introduced by Fillmore (1982) as a theory of meaning in which words are understood relative to structured background knowledge called frames. A frame is a conceptual scene (e.g., *Commercial\_transaction*, *Motion*) that specifies the key participants, relations, and

<sup>1</sup>All the codes for parsing, preprocessing and experimenting are included in <https://github.com/phantaatbach/STARSEM26>.

<sup>2</sup>We could only test this method on the English dataset due to the lack of available and reliable Frame Semantics resources for other languages. This is justified in section 3.

<sup>3</sup>While we acknowledge the growing field of explainable AI and its diverse range of techniques (e.g., Integrated Gradients (Sundararajan et al., 2017), LIME (Ribeiro et al., 2016) or SHAP (Lundberg and Lee, 2017), see surveys in Arrieta et al. (2019); Madsen et al. (2022)), these methods are often post-hoc, add methodological and computational overhead, and may not yield explanations that are straightforward to inspect linguistically. This makes our approach particularly convenient for linguistic and humanities research, where direct, feature-level inspection is desired.

typical inferences involved. The meaning of a lexical unit is not just a standalone definition, but a particular way of evoking a frame and profiling parts of it. For example, *buy*, *sell*, *pay*, and *cost* all relate to the same *commercial* frame, but each highlights different roles and perspectives. Frames come with frame elements that represent the core components of the scene (e.g., *buyer*, *seller*, *goods*, *money*). The theory is operationalised in resources such as FrameNet (Baker, 2014), which catalogs frames, frame elements, and the lexical units that evoke them, along with annotated examples. Because of its structured representation of ‘who did what to whom’ at the level of conceptual situations, Frame Semantics representations have supported applications in semantic parsing, information extraction, and linguistic analysis (e.g., Law, 2019; Li et al., 2019; Marzinotto et al., 2018; Robin et al., 2023; Vais, 2022).

Some practical weaknesses of Frame Semantics are:

- The substantial time and manual work needed for building corpora.
- The lack of coverage for different languages as frame inventories are built from particular languages and cultural/institutional settings, although efforts have been made for cross-lingual applicability (Boas et al., 2025).
- The lack of coverage for historical and out-of-domain data as many senses, lexical items, constructions are missing from the inventory (Hartmann et al., 2017; Palmer and Sporleder, 2010).

These may explain why it has mainly been used in specific linguistics case studies (Blank, 1999, 2003; Georgakopoulos, 2018; Koch, 1999a, 2008, 1999b; Law, 2019; Vais, 2022) and not in unsupervised computational approaches to LSC detection. As far as we know, our work is the first one to employ Frame Semantics for this task.

### 3 Data and tasks

Following SemEval 2020 Task 1 on Unsupervised Lexical Semantic Change Detection (Schlechtweg et al., 2020), we formulate the problem as either binary classification or graded ranking, corresponding to subtasks 1 and 2. In subtask 1, given a set of target words, the system predicts whether each word gained or lost at least one sense between two

time periods. In subtask 2, the system ranks target words by their degrees of semantic change. Annotating semantic change is challenging because it requires assessing meaning differences across many usage instances from each period. A widely used protocol is DUREl (Schlechtweg et al., 2018), where annotators compare pairs of sentences containing a target word and judge whether the instances are used in the same sense or in different senses. These pairwise judgments are then aggregated into a final change score, either binary or continuous. This procedure was used to create the SemEval 2020 Task 1 dataset, the most widely used one to compare different LSC detection systems. It covers four languages: English, German, Latin, and Swedish, with manual annotations for both subtasks, and provides a diachronic corpus split into two time periods for each language. Using this benchmark enables direct comparison of our approach with SOTA systems under a shared evaluation setup.

However, our approach requires high-quality frame-semantic parsers, which are not currently available for all languages. While English benefits from extensive FrameNet resources and several mature automatic frame-semantic parsers, we were unable to find comparably reliable off-the-shelf parsers for the other 3 languages. For English, we used the pretrained Frame Semantics Transformer model (Chanin, 2023), which is trained on FrameNet 1.7 and is reported to achieve SOTA performance. It is released as a Python library, making it convenient for using and reproducing the results. For German, although there are several FrameNet-like resources, they were all relatively dated (Boas, 2002; Burchardt et al., 2009, 2006; Erk and Padó, 2006). The most recent resource and parser we have found is PFN-DE (Bick, 2022). However, it is only accessible through a hosted interface that only accepts text files up to 10,000 characters, and processing larger volumes requires contacting the provider for a paid processing agreement (VISL, 2026). This makes it impractical to parse 140 million token historical German corpora (Schlechtweg et al., 2020) at scale compared to a locally runnable, openly distributed pipeline, and raises concern about transparency and reproducibility. Moreover, it has only been evaluated on modern Twitter data, making its reliability on historical German uncertain. For Swedish, Svenskt frasnät 2.0 (SweFN) Dataset (Språkbanken Text, 2024) provides a FrameNet-style resource. The latest at-

Period	Tokens	Types	Corpus	Target words
C1 (1810 – 1860)	6.5M	Raw, Lemma	CCOHA	37
C2 (1960 – 2010)	6.7M	Raw, Lemma	CCOHA	

Table 1: Descriptions of the English dataset in SemEval 2020 Task 1.

tempt to build a pretrained Swedish frame parser with SweFN used the same architecture as the English ones mentioned above, but with substantially lower performance (Dannélls and Virk, 2025; Dannélls et al., 2024). In a small sanity check with 30 random, diverse and basic sentences, we also observed 13 questionable frame predictions, such as *Cutting for Jag såg en film igår* ‘I saw a movie yesterday’ (see the results in table 4 in the Appendix). The low F1 scores, coupled with its failure in such simple examples, raise concern about the reliability of the model. For Latin, we did not find any public frame-semantic resources or pretrained parsers.

Because of the lack of available and quality frame-semantic parsers for German, Swedish and Latin, we restrict the experiment to English where strong pretrained parsers exist. We compensate for the lack of multilingual quantitative experiments with a detailed qualitative analysis. A quick description of the English data is presented in Table 1. Note that although we evaluate only English, frame semantics is in principle language-agnostic (Fillmore, 1982) so the approach should generalise once suitable frame resources and parsers exist.

## 4 Method

### 4.1 FrameNet parsing

Because we only need sentences containing the target lemma, we first split both the lemmatised and raw data into lemma-specific subcorpora. Since the lemmatised and raw files contain the same number of sentences in the same order, we do this by scanning the lemmatised file to select sentences that include the target lemma, then retrieving the corresponding sentences from the raw (token) file by index. We then frame-parsed them using the Frame transformer model (Chanin, 2023). The parser is available in two variants, BASE and SMALL, with similar reported performances. It has also been noted that the models can occasionally identify a frame but return empty outputs. We therefore use BASE by default and fall back to SMALL when it fails. If both models fail, we skip the sentence. Across 49,416 sentences, we triggered the fallback

for 2,413 lemmatised sentences (4.9%) and 1,052 raw sentences (2.1%). Amongst these, both models failed (and the sentences were skipped) for only 805 lemmatised sentences (1.6%) and 336 raw sentences (0.6%).

One technical limitation of the parser is that it returns the frame elements as raw strings without locations. This makes it impossible to locate the target lemma in the parsed token data (one target lemma may appear in different surface forms). We therefore use lemmatised corpora rather than raw corpora. To assess the impact of doing so, we compared (within each target word and time period) the frame distributions derived from the lemmatised and raw corpora using Jensen–Shannon Divergence (Menéndez et al., 1997). The resulting JSD values (in the table 5 in the Appendix) are uniformly small (all  $< 0.15$ , most  $< 0.10$ ), indicating that using lemmatised corpora introduces only minor changes to the induced frame distributions.

### 4.2 Basic procedure

A frame-parsed sentence consists of a sentence, one or more frame instances, the lexical item that evokes each frame, and the corresponding frame elements. For example, the sentence *Brussels Airline buy 100 plane\_nn from Boeing* contains a COMMERCE\_BUY frame, with *buy* as the frame-trigger and the target lemma *plane\_nn* inside the GOODS frame-element.

To obtain the frame distributions, we collected all annotated frame instances in which the target lemma occurs either as the frame-evoking element or within the lexical span of one of the frame elements.

- Each frame instance was counted at most once. Even if the target lemma appeared both as the frame trigger and inside a frame element, or appeared in multiple frame-element spans of the same frame instance, that frame instance was still counted only once. This de-duplication was applied only within individual frame instances.
- If the target lemma participated in multiple

frame instances in the same sentence, all such frame instances were retained and counted as separate occurrences.

We then experiment with 2 types of frame distributions<sup>4</sup>: frame–elements only (FE) and frame–triggers + frame–elements (FTFE). To quantify the distributional changes across periods, we use Jensen–Shannon divergence (JSD) (Menéndez et al., 1997). JSD is a symmetric, smoothed variant of Kullback–Leibler Divergence (Kullback and Leibler, 1951), and with log base 2 it is bounded between 0 and 1, making scores easy to interpret and comparable across targets. In addition, JSD can be decomposed into per-item contributions, allowing fine-grained attribution of which items (in this case, frames) drive the observed change. Prior work on grammatical profiles (Kutuzov et al., 2021; Ryzhova et al., 2021) quantified changes with cosine distance between time-specific vectors, but cosine compares vector geometry (angles) rather than directly comparing probability distributions.

## 5 Results

We evaluate our method on both subtasks of SemEval 2020 Task 1. As described in Section 3, subtask 1 is a binary classification task evaluated by accuracy, while subtask 2 is a ranking task evaluated by Spearman’s rank correlation between the JSDs and the GOLD SCORES. Our primary focus is Subtask 2. For Subtask 1, we classify a word as CHANGED if its JSD  $\geq 0.5$ , else UNCHANGED.

### 5.1 Subtask 2

The results for subtask 2 are reported in table 2. Using frame-elements only, we were able to achieve a spearman correlation of 0.249. This is already much higher than the Count Baseline (0.022) and surpassing many static and contextualised embedding systems. Using the combined frame-triggers and frame-elements boosts the score to 0.306, making our system amongst the top 10 systems for English. This is very impressive given the simplicity of the approach.

### 5.2 Subtask 1

The results for subtask 1 are presented in table 3. As the experiment set-up is fully unsupervised, dif-

<sup>4</sup>We intentionally excluded the frame – trigger only distribution. This is because some lemmas only appear in frame elements and not evoke any frame, making using frame – triggers only yield sparse or even empty distributions for some targets, causing JSD to be unstable or even undefined.

Team	EN score	System
XL-Lexeme	0.757	token
NLPCR	0.436	token
UG_Student_Intern	0.422	type
cs2020	0.375	type
UWB	0.367	type
Discovery_Team	0.361	ens.
Jiaxin & Jinan	0.325	type
UCD	0.307	graph
<b>FrameNet-FTFE</b>	<b>0.306</b>	–
IMS	0.301	type
Life-Language	0.299	type
Syntactic Dependency	0.277	–
Entity	0.250	type
<b>FrameNet-FE</b>	<b>0.249</b>	–
RPI-Trust	0.228	type
Grammatical Profiling	0.218	–
Random	0.211	type
Skurt	0.209	token
RIJP	0.157	type
UiO-UvA	0.136	token
UoB	0.105	topic
cbk	0.059	token
NLP@IDSIA	0.028	token
<b>Count Bas.</b>	0.022	–
JCT	0.014	type
TUE	-0.155	token
DCC	-0.217	type
<b>Freq. Bas.</b>	-0.217	–
<b>Maj. Bas.</b>	–	–

Table 2: Performances of our models in comparison with other systems in graded change detection (SemEval 2020 Subtask 2), Spearman rank correlation coefficients. FTFE = Frame-Trigger-Frame-Element, FE = Frame Element.

ferent teams come up with different strategies to classify the changing status of the target lemmas. We based our binary classification score assignment on the 0.5 cutoff point of JSD. The threshold of 0.5 was chosen empirically, based on our prior experience using JSD on linguistic data. It acts as a high-pass filter for semantic salience: it highlights substantial structural shifts while suppressing minor fluctuations that often arise from sparse distributions. We avoid any supervised finetuning for selecting the optimal JSD threshold to preserve the fully unsupervised spirit of the task and comparability with other unsupervised approaches. Both of our systems perform equally well, with an accuracy

Team	EN score	System type
NLPCR	0.730	token
Life-Language	0.703	type
Entity	0.676	type
Jiaxin & Jinan	0.649	type
RPI-Trust	0.649	type
DCC	0.649	type
JCT	0.649	type
Syntactic Dependency	0.649	–
<b>FrameNet-FTFE</b>	<b>0.622</b>	–
<b>FrameNet-FE</b>	<b>0.622</b>	–
UWB	0.622	type
Grammatical Profiling	0.622	–
NLP@IDSIA	0.622	token
UCD	0.622	graph
<b>Count Bas.</b>	0.595	–
cs2020	0.595	token
UG_Student_Intern	0.568	type
Skurt	0.568	token
Discovery_Team	0.568	ens.
TUE	0.568	token
<b>Maj. Bas.</b>	0.568	–
cbk	0.568	token
UoB	0.568	topic
IMS	0.541	type
UiO-UvA	0.541	token
RIJP	0.541	type
Random	0.486	type
<b>Freq. Bas.</b>	0.432	–

Table 3: Performances of our models in comparison with other systems in binary change detection (SemEval 2020 Subtask 1), accuracy. Note that in this paper we mostly focus on ranking (Subtask 2). FTFE = Frame-Trigger-Frame-Element, FE = Frame-Element.

of 0.622, on par with different word embedding systems.

## 6 Qualitative analysis

A key strength of our system is its interpretability. JSD decomposition reveals which distributional shifts drive the change, and users can inspect the corresponding contexts to locate the underlying semantic change. In this section, we qualitatively analyse the predictions of our better system (FTFE) to investigate what has been captured.

We follow the same strategy as in our previous study (Tat et al., 2026), where we divided the target lemmas into 5 groups: True Positive (TP), True Negative (TN), False Positive (FP), False Negative

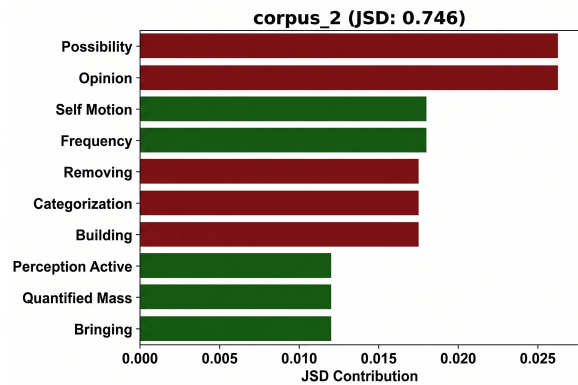


Figure 2: The JSD contributions of different frames of the English noun *prop*. Green bars indicate an increase in relative frequency; red bars indicate a decrease.

(FN), MID. We define TP as lemmas that fall in the top 33% of both the Subtask 2 Gold Scores and our system’s predicted rankings, and TN as those in the bottom 33% of both. FP are lemmas in the bottom 33% of the gold ranking but the top 33% of our ranking, and FN are lemmas in the top 33% of the gold ranking but the bottom 33% of our ranking. All remaining lemmas are assigned to a MID group. Our qualitative analysis focuses on TP, TN, FP, and FN, as these categories most clearly expose the method’s strengths and limitations. Due to space constraints, we present one illustrative (most extreme) example per group. The lemmas of each group are listed in Table 6 in the Appendix.

**True positives** Upon qualitative analysis, semantic changes do indeed link to changes in frame distributions. The English noun *prop* develops a new THEATRE OBJECT meaning whereas in the earlier period the SUPPORT meaning (either literal or figurative) was dominant. In the first period, *prop* were mostly linked to such frames as *Possibility* (e.g., the possibility of finding support) or *Opinion* (e.g, thinking about the pillar/bread-winner of the family, or opposing the election of the prop/council of the city) or *Building*. In the second period, those frames decrease and there was an increase in *Self motion* frames, where *prop* acts as the mover (e.g., ... *props fly into scene*) or looking for the *prop* is the purpose of the mover (e.g., *he clambered ... looking for perfect props*) or *Quantified mass* where *prop* was quantified by *few*, *many* or *Bringing* where *prop* is the object being brought. This is illustrated in figure 2.

**True negatives** The JSD scores of TNs are really low, suggesting that there are minor changes in the




Figure 3: The JSD contributions of different frames of the English noun *tree*. Green bars indicate an increase in relative frequency; red bars indicate a decrease.




Figure 4: The JSD contributions of different frames of the English noun *quilt*. Green bars indicate an increase in relative frequency; red bars indicate a decrease.

frames they appear in. The major shift of the English noun *tree* (see Figure 3) is the fact the phrase *old tree* appears less (C1: 76 times, C2: 9 times). Other shifts are minor (hence the small overall JSD) and simply contextual (e.g., the decrease in *trees* being destroyed), not related to any real semantic change.

**False positives** We only found one FP (*quilt\_nm*). Our method detects large shifts in its frame distributions, yet it is ranked low in the gold Subtask 2 ranking. This is because of the large shift in the contexts (which do not correspond to any semantic change) in which *quilt* appears in. In the second period, *quilt* was the object of many activities, such as *lean into the quilt*, *wring the quilt*, *lay the folded quilt*, etc. This results in a large increase of such frames as *Posture*, *Manipulation*, *Placing*, *Spatial contact*, *Causation*, which drives the overall JSD up (see figure 4).

**False negative** FNs are target words whose frame JSD scores underestimate the degree of change rel-




Figure 5: The JSD contributions of different frames of the English noun *ball*. Green bars indicate an increase in relative frequency; red bars indicate a decrease.

ative to the gold ranking. For example, our system does capture the shift for the English noun *ball* from a usage dominated by *SOCIAL EVENT* to one dominated by *SPHERICAL OBJECT*, reflected in decreases in frames such as *Social event* and *Building subpart* (e.g., *ball room*) and increases in frames such as *Cause motion* and *Manipulation* (see figure 5). We hypothesise that, for FN cases, these distributional shifts are just simply not large enough to place the target high in the ranked list. Nevertheless, the method remains useful for qualitative analysis, since it highlights the relevant frame shifts and supports interpretation of semantic change through the lens of Frame Semantics.

There are two common problems we observe during the qualitative analysis. First, some frames are highly general and carry little specific semantic content. For instance, in the *ball* case, the rise of the *Statement* frame largely reflects the frequent occurrences of *say* in contexts that mention *ball* (whether referring to the social event or the object), rather than a substantive meaning shift. Second, using lemmatised corpora can occasionally disrupt grammatical agreement or sentence boundaries, which may confuse the parser and lead to incorrect frame assignments (this effect appears minor overall, as shown by the low JSD between the raw and lemmatised corpora’s frame inventories).

## 7 Conclusion

This paper, along with other linguistically motivated case studies on SemEval 2020 Task 1 (e.g., Tat et al., 2026; Kutuzov et al., 2021; Ryzhova et al., 2021; Tang et al., 2023) demonstrates that explicit linguistic evidence is sufficient to build an effective lexical semantic change detection system.

Our method achieves competitive performance on English and can outperform many distributional baselines on both subtasks. Although it does not match the current SOTA systems which are based on large pretrained language models and can encode a broader range of linguistic and contextual signals, maximising benchmark scores has never been our primary goal. Rather, we aim to highlight and thus bring more attention to theory-driven, interpretable approaches to lexical semantic change detection.

To further test and refine this approach, future work should:

- Develop Frame inventories and more reliable Frame parsers for other languages
- Work around the problem of not being able to locate the locations of frame elements.

More broadly, we view this approach and some previous ones (e.g., Kutuzov et al., 2021; Tat et al., 2026; Ryzhova et al., 2021) as instances of a general framework for semantic change analysis: identify interpretable, theory-driven dimensions of a target’s distribution (here, frames) and quantify change within each dimension. This framework could (and should) be used as a complementary tool to interpret and validate the changes detected by neural networks.

## Limitations

We highlight two primary limitations and practical considerations of our proposed method.

First, frame inventory and frame parsers availability are a hard constraint, making this method inapplicable to languages lacking robust frame parsers. Even for English, the frame parser is still not perfect. It cannot output the token offsets for frame elements, meaning we must use the lemmatised version of the corpora, which can (minorly) affect the performance of the approach.

Second, our qualitative analysis, while supporting the interpretability claim, simultaneously reveals one inherent limitation of distributional frame tracking. The method can overestimate changes when large shifts in a target’s contexts do not correspond to genuine semantic change. Conversely, it can underestimate change when genuine semantic innovations are not sufficiently reflected in frame distributional shifts.

## Ethical statement

The data and all codes used in this study are publicly available. No human data was used, and no human participant was involved in the experiment.

## Acknowledgments

All ideas and core content were developed by the authors then later polished and optimised using ChatGPT and Gemini. We reviewed all AI-generated content thoroughly and take full responsibility for its accuracy.

This research is part of CASCADE (101119511), a Horizon Europe MSCA Doctoral Network funded by the European Commission and UKRI.

We would like to thank the anonymous reviewers for their detailed and informative comments and suggestions.

## References

- Josep Alba-Salas. 2007. [On the life and death of a collocation: A corpus-based diachronic study of dar miedo/hacer miedo-type structures in Spanish](#). *Diachronica*, 24(2):207–252. Type: Journal Article.
- Laurence Anthony. 2014. [AntConc \(Version 3.4.3\)](#)[Computer Software]. Tokyo, Japan: Waseda University.
- Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bannetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2019. [Explainable Artificial Intelligence \(XAI\): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI](#). *arXiv preprint: 1910.10045*.
- Collin F. Baker. 2014. [FrameNet: A Knowledge Base for Natural Language Processing](#). In *Proceedings of Frame Semantics in NLP: A Workshop in Honor of Chuck Fillmore (1929-2014)*, pages 1–5, Baltimore, MD, USA. Association for Computational Linguistics.
- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. 2014. [Don’t count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 238–247, Baltimore, Maryland. Association for Computational Linguistics.
- Eckhard Bick. 2022. [A Framenet and Frame Annotator for German Social Media](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3942–3949, Marseille, France. European Language Resources Association.

- Andreas Blank. 1999. [Why do new meanings occur? A cognitive typology of the motivations for lexical semantic change](#). In Andreas Blank and Peter Koch, editors, *Historical Semantics and Cognition*, pages 61–90. De Gruyter Mouton, Berlin, Boston.
- Andreas Blank. 2003. [Words and concepts in time: Towards diachronic cognitive onomasiology](#). In Regine Eckardt, Klaus von Heusinger, and Christoph Schwarze, editors, *Diachronic Semantics from Different Points of View*, pages 37–66. De Gruyter Mouton, Berlin, New York.
- Hans C. Boas. 2002. [Bilingual FrameNet Dictionaries for Machine Translation](#). In *Proceedings of the Third International Conference on Language Resources and Evaluation (LREC'02)*, Las Palmas, Canary Islands - Spain. European Language Resources Association (ELRA).
- Hans C Boas, Josef Ruppenhofer, and Collin F Baker. 2025. [FrameNet at 25: Results and Applications](#). *International Journal of Lexicography*, 38(2):159–189. [\\_eprint: https://academic.oup.com/ijl/article-pdf/38/2/159/63310662/ecaf011.pdf](https://academic.oup.com/ijl/article-pdf/38/2/159/63310662/ecaf011.pdf).
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching Word Vectors with Subword Information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146. Place: Cambridge, MA.
- Vaclav Brezina, Pierre Weill-Tessier, and Tony McEnery. 2020. [#LancsBox \(version 5.x\)](#).
- John A. Bullinaria and Joseph P. Levy. 2007. [Extracting semantic representations from word co-occurrence statistics: A computational study](#). *Behavior Research Methods*, 39(3):510–526.
- Aljoscha Burchardt, Katrin Erk, Anette Frank, Andrea Kowalski, Sebastian Padó, and Manfred Pinkal. 2006. [The SALSA Corpus: a German Corpus Resource for Lexical Semantics](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Aljoscha Burchardt, Katrin Erk, Anette Frank, Andrea Kowalski, Sebastian Padó, and Manfred Pinkal. 2009. [8. Using FrameNet for the semantic analysis of German: Annotation, representation, and automation](#). In Hans C. Boas, editor, *Methods and Applications*, pages 209–244. De Gruyter Mouton, Berlin, New York.
- 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.
- David Chanin. 2023. [Open-source Frame Semantic Parsing](#). *arXiv preprint arXiv:2303.12788*.
- Kenneth Ward Church and Patrick Hanks. 1990. [Word Association Norms, Mutual Information, and Lexicography](#). *Computational Linguistics*, 16(1):22–29.
- Dana Dannélls, Richard Johansson, and Lucy Yang Buhr. 2024. [Transformer-based Swedish Semantic Role Labeling through Transfer Learning](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16762–16769, Torino, Italia. ELRA and ICCL.
- Dana Dannélls and Shafqat Virk. 2025. [A Frame-Semantic Parsing Plugin for Swedish Research Infrastructure](#). In *Proceedings of the 2nd Huminfra Conference (HiC 2025)*, volume 60 of *NEALT Proceedings Series*, pages 8–12.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *North American Chapter of the Association for Computational Linguistics*.
- Katrin Erk and Sebastian Padó. 2006. [Shalmaneser - A Toolchain For Shallow Semantic Parsing](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Charles J. Fillmore. 1982. [Frame semantics](#). In *Linguistics in the Morning Calm: Selected Papers from SICOL-1981*, pages 111–137, Seoul. Hanshin Publishing Co.
- John Rupert Firth. 1957. [A synopsis of linguistic theory 1930–1955](#). In *Studies in Linguistic Analysis*, pages 1–32. Basil Blackwell, Oxford.
- Marcos Garcia and Marcos García Salido. 2019. [A Method to Automatically Identify Diachronic Variation in Collocations](#). In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 71–80, Florence, Italy. Association for Computational Linguistics.
- D. Geeraerts. 1983. [Prototype Theory and Diachronic Semantics. A Case Study](#). *Indogermanische Forschungen (1983)*, 88(1983):1–32.
- Dirk Geeraerts. 2010. [Theories of lexical semantics](#). Oxford University Press, Oxford, New York.
- Dirk Geeraerts, Dirk Speelman, Kris Heylen, Mariana Montes, Stefano De Pascale, Karlien Franco, and Michael Lang. 2023. [Lexical Variation and Change: A Distributional Semantic Approach](#), 1 edition. Oxford University Press Oxford.
- Thanasis Georgakopoulos. 2018. [A frame-based approach to the source-goal asymmetry](#). *Constructions and Frames*, 10(1):61–97. Type: Journal Article.

- 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.
- Stefan Th. Gries and Dagmar Divjak. 2009. [Behavioral profiles: A corpus-based approach to cognitive semantic analysis](#). In Vyvyan Evans and Stéphanie Pourcel, editors, *Human Cognitive Processing*, volume 24, pages 57–75. John Benjamins Publishing Company, Amsterdam.
- Kristina Gulordava and Marco Baroni. 2011. [A distributional similarity approach to the detection of semantic change in the Google Books Ngram corpus](#). In *Proceedings of the GEMS 2011 Workshop on GEometrical Models of Natural Language Semantics*, pages 67–71, Edinburgh, UK. Association for Computational Linguistics.
- Silvana Hartmann, Iliia Kuznetsov, Teresa Martin, and Iryna Gurevych. 2017. [Out-of-domain FrameNet Semantic Role Labeling](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 471–482, Valencia, Spain. Association for Computational Linguistics.
- Ignacio Iacobacci, Mohammad Taher Pilehvar, and Roberto Navigli. 2016. [Embeddings for Word Sense Disambiguation: An Evaluation Study](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 897–907, Berlin, Germany. Association for Computational Linguistics.
- Marlies Jansegers and Stefan Th. Gries. 2020. [Towards a dynamic behavioral profile: A diachronic study of polysemous sentir in Spanish](#). *Corpus Linguistics and Linguistic Theory*, 16(1):145–187.
- Bryan Jurish. 2015. DiaCollo: On the trail of diachronic collocations.
- Adam Kilgarriff, Vít Baisa, Jan Bušta, Miloš Jakubíček, Vojtěch Kovář, Jan Michelfeit, Pavel Rychlý, and Vít Suchomel. 2014. The Sketch Engine: ten years on. *Lexicography*, 1:7–36.
- Yoon Kim, Yi-I Chiu, Kentaro Hanaki, Darshan Hegde, and Slav Petrov. 2014. [Temporal Analysis of Language through Neural Language Models](#). In *Proceedings of the ACL 2014 Workshop on Language Technologies and Computational Social Science*, pages 61–65, Baltimore, MD, USA. Association for Computational Linguistics.
- Peter Koch. 1999a. [Cognitive aspects of semantic change and polysemy: the semantic space HAVE/BE](#). In Andreas Blank and Peter Koch, editors, *Historical Semantics and Cognition*, pages 279–306. De Gruyter Mouton, Berlin, Boston.
- Peter Koch. 1999b. [Frame and Cognition: On the Cognitive Bases of Metonymy and Certain Types of Word Formation](#). In Klaus-Uwe Panther and Günter Radden, editors, *Metonymy in Language and Thought*, pages 139–168. John Benjamins Publishing Company.
- Peter Koch. 2008. [Cognitive onomasiology and lexical change: Around the eye](#). In Martine Vanhove, editor, *Towards a typology of lexical semantic associations*, pages 107–137. John Benjamins Publishing Company.
- Vivek Kulkarni, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. [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, CHE. International World Wide Web Conferences Steering Committee. Event-place: Florence, Italy.
- S. Kullback and R. A. Leibler. 1951. [On Information and Sufficiency](#). *The Annals of Mathematical Statistics*, 22(1):79–86.
- Andrey Kutuzov, Lidia Pivovarova, and Mario Giulianelli. 2021. [Grammatical Profiling for Semantic Change Detection](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 423–434, Online. 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*, pages 1384–1397, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- George Lakoff and Mark Johnson. 1981. *Metaphors we live by*. University of Chicago Press, Chicago.
- James Law. 2019. [Diachronic frame analysis](#). *Constructions and Frames*, 11(1):43–78. Type: Journal Article.
- Alessandro Lenci, Magnus Sahlgren, Patrick Jeuniaux, Amaru Cuba Gyllensten, and Martina Miliani. 2022. [A comparative evaluation and analysis of three generations of Distributional Semantic Models](#). *Language Resources and Evaluation*, 56(4):1269–1313.
- Omer Levy and Yoav Goldberg. 2014. [Dependency-Based Word Embeddings](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 302–308, Baltimore, Maryland. Association for Computational Linguistics.
- Wei Li, Dezhi Cheng, Lei He, Yuanzhuo Wang, and Xiaolong Jin. 2019. [Joint Event Extraction Based on Hierarchical Event Schemas From FrameNet](#). *IEEE Access*, 7:25001–25015.
- Scott M. Lundberg and Su-In Lee. 2017. [A Unified Approach to Interpreting Model Predictions](#). In *Neural Information Processing Systems*.

- Andreas Madsen, Siva Reddy, and Sarath Chandar. 2022. [Post-hoc Interpretability for Neural NLP: A Survey](#). *ACM Computing Surveys*, 55(8):1–42.
- Gabriel Marzinotto, Jeremy Auguste, Frederic Bechet, Geraldine Damnati, and Alexis Nasr. 2018. [Semantic Frame Parsing for Information Extraction : the CALOR corpus](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Tony McEnery, Vaclav Brezina, and Helen Baker. 2019. [Usage Fluctuation Analysis: A new way of analysing shifts in historical discourse](#). *International Journal of Corpus Linguistics*, pages 413–444.
- M.L. Menéndez, J.A. Pardo, L. Pardo, and M.C. Pardo. 1997. [The Jensen-Shannon divergence](#). *Journal of the Franklin Institute*, 334(2):307–318.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). *arXiv preprint*. ArXiv:1301.3781 [cs].
- Syrielle Montariol, Matej Martinc, and Lidia Pivovarova. 2021. [Scalable and Interpretable Semantic Change Detection](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4642–4652, Online. Association for Computational Linguistics.
- Alexis Palmer and Caroline Sporleder. 2010. [Evaluating FrameNet-style semantic parsing: the role of coverage gaps in FrameNet](#). In *Coling 2010: Posters*, pages 928–936, Beijing, China. Coling 2010 Organizing Committee.
- Jeffrey Pennington, Richard Socher, and Christopher 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, Doha, Qatar. Association for Computational Linguistics.
- 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, New Orleans, Louisiana. Association for Computational Linguistics.
- Daniela Pettersson-Traba. 2016. [Analyzing the behavioral profiles of sets of near-synonyms in American English from a diachronic perspective](#).
- Daniela Pettersson-Traba. 2021. [Measuring semantic distance across time](#). *Journal of Research Design and Statistics in Linguistics and Communication Science*, 6(2):138–165. [\\_eprint: https://doi.org/10.1558/jrds.40245](https://doi.org/10.1558/jrds.40245).
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [WiC: the Word-in-Context Dataset for Evaluating Context-Sensitive Meaning Representations](#). 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 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Marco Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. [“Why Should I Trust You?”: Explaining the Predictions of Any Classifier](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 97–101, San Diego, California. Association for Computational Linguistics.
- Cécile Robin, Atharva Kulkarni, and Paul Buitelaar. 2023. [Identifying FrameNet Lexical Semantic Structures for Knowledge Graph Extraction from Financial Customer Interactions](#). In *Proceedings of the 12th Global Wordnet Conference*, pages 91–100, University of the Basque Country, Donostia - San Sebastian, Basque Country. Global Wordnet Association.
- Anastasiia Ryzhova, Daria Ryzhova, and Ilya Sochenkov. 2021. [Detection of Semantic Changes in Russian Nouns with Distributional Models and Grammatical Features](#).
- Dominik Schlechtweg, Barbara McGillivray, Simon Hengchen, Haim Dubossarsky, and Nina Tahmasebi. 2020. [SemEval-2020 Task 1: Unsupervised Lexical Semantic Change Detection](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1–23, Barcelona (online). International Committee for Computational Linguistics.
- 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, Volume 2 (Short Papers)*, pages 169–174, New Orleans, Louisiana. Association for Computational Linguistics.
- Hinrich Schütze. 1998. [Automatic Word Sense Discrimination](#). *Computational Linguistics*, 24(1):97–123. Place: Cambridge, MA.
- John Sinclair. 1991. [Corpus, Concordance, Collocation](#).
- Tamar Sovran. 2004. [Polysemy and Semantic Frames - A diachronic study of "Truth" in Hebrew](#). 26(1):89–100.

- Språkbanken Text. 2024. [Svenskt frasnät 2.0 \(SweFN\)](#).
- Karen Sullivan. 2007. Metaphoric Extension and Invited Inferencing in Semantic Change. *Culture, Language and Representation*, Special Issue: Metaphor and Discourse:257–274.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. [Axiomatic Attribution for Deep Networks](#). In *International Conference on Machine Learning*.
- Eve Sweetser. 1990. Semantic structure and semantic change: English perception-verbs in an Indo-European context. In *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*, Cambridge Studies in Linguistics, pages 23–48. Cambridge University Press.
- Nina Tahmasebi, Lars Borin, and Adam Jatowt. 2021. [Survey of computational approaches to lexical semantic change detection](#). In *Computational approaches to semantic change*. Zenodo.
- Nina Tahmasebi and Haim Dubossarsky. 2023. [Computational modeling of semantic change](#). [\\_eprint: 2304.06337](#).
- Xiaohang Tang, Yi Zhou, Taichi Aida, Procheta Sen, and Danushka Bollegala. 2023. [Can Word Sense Distribution Detect Semantic Changes of Words?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3575–3590, Singapore. Association for Computational Linguistics.
- Bach Phan Tat, Kris Heylen, Dirk Geeraerts, Stefano De Pascale, and Dirk Speelman. 2026. [Transparent semantic change detection with dependency-based profiles](#). In *The Proceedings for the 6th International Workshop on Computational Approaches to Language Change (LChange'26)*, pages 97–109, Rabat, Morocco. Association for Computational Linguistics.
- Christine Marie Vais. 2022. *A Diachronic Application of Frame Semantics to Meaning Change in Prepositions: The Case of German mit*. Doctoral dissertation, University of California, Berkeley, Berkeley, CA.
- VISL. 2026. [VISL - Upload interface](#).
- David Yenicelik, Florian Schmidt, and Yannic Kilcher. 2020. [How does BERT capture semantics? A closer look at polysemous words](#). In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 156–162, Online. Association for Computational Linguistics.

## Appendix

### A Testing the Swedish model on 30 sentences

Swedish	English translation	Predicted verbal frames - SWE	Predicted verbal frames - ENG
Jag såg en film igår	I saw a movie yesterday	<b>Cutting</b>	Perception_experience
Vi äger ett hus	We own a house	<b>Cause_to_fragment</b>	Possession
Vi lånar pengar	We borrow money	Borrowing	Getting
Mamma skickar ett brev	Mom sends a letter	Passing	Sending
Barnet dricker mjölk	The child drinks milk	Ingestion	Ingestion
Hon säger sanningen	She tells the truth	<b>Deciding</b>	Telling
Jag vet svaret	I know the answer	<b>Health_status</b>	Awareness
Bilen stannar här	The car stops here	Halt	Halt
Han öppnar fönstret	He opens the window	<b>Closure</b>	Closure
Vi går till parken	We go to the park	Motion	Motion
Tåget anländer nu	The train arrives now	Arriving	Arriving
Boken ligger på hyllan	The book is on the shelf	<b>Posture</b>	No verbal frame. Only Text triggered by <i>book</i>
Jag tänker på dig	I think of you	Cogitation	Awareness
Han älskar sin hund	He loves his dog	Experiencer_focus	Experiencer_focus
Hon förstår frågan	She understands the question	<b>Frugality</b>	Awareness
Vi glömmer namnet	We forget the name	Memory	Remembering_information
De pratar om vädret	They talk about the weather	Chatting	Chatting
Läraren svarar eleven	The teacher answers the student	<b>Rising_to_a_challenge</b>	Education_teaching
Pappa lovar en present	Dad promises a present	Commitment	<b>Cause_to_perceive</b>
Hon ringer sin vän	She calls her friend	<b>Cause_to_make_noise</b>	Referring_by_name
Isen smälter snabbt	Ice melts quickly	Change_of_phase	Change_of_phase
Solen skiner idag	The sun shines today	Light_movement	Light_movement
Han slår igen dörren	He slams the door	<b>Fall_for</b>	Cause_impact
Han skriver ett mejl	He writes an email	Text_creation	Text_creation
Mötet börjar klockan nio	The meeting starts at nine o'clock	Process_start	Process_start
Jag slår på TV:n	I turn on the TV	<b>Make_noise</b>	Change_operational_state
Det regnar ute	It is raining outside	Precipitation	Precipitation
Barnet sover djupt	The child sleeps deeply	Sleep	Sleep
Hon kommer på svaret	She comes up with the answer	<b>Satisfying</b>	Coming_up_with
Hon ger upp	She gives up	<b>Excreting</b>	EMPTY FRAME

Table 4: Testing the Swedish models on 30 random and basic sentences, amongst which 13 frames were questionable (in bold). Compare with the results from the English BASE model..

### B JSD between frame distributions of raw and lemmatised English corpora

<b>Target lemma</b>	<b>JSD period 1 (token vs lemma)</b>	<b>JSD period 2 (token vs lemma)</b>
attack_nn	0.053278	0.036771
bag_nn	0.093656	0.049945
ball_nn	0.077025	0.076959
bit_nn	0.073138	0.048374
chairman_nn	0.086464	0.042575
circle_vb	0.118216	0.127109
contemplation_nn	0.08533	0.131131
donkey_nn	0.118224	0.097751
edge_nn	0.069459	0.048838
face_nn	0.017471	0.017577
fiction_nn	0.085575	0.065596
gas_nn	0.142413	0.072828
graft_nn	0.126597	0.109724
head_nn	0.017066	0.018624
land_nn	0.02214	0.027298
lane_nn	0.067174	0.065796
lass_nn	0.106964	0.134866
multitude_nn	0.065254	0.130617
ounce_nn	0.122896	0.125564
part_nn	0.03803	0.05156
pin_vb	0.109057	0.102104
plane_nn	0.142745	0.074554
player_nn	0.146468	0.071431
prop_nn	0.134356	0.14812
quilt_nn	0.088169	0.090321
rag_nn	0.09174	0.117761
record_nn	0.068515	0.047556
relationship_nn	0.122012	0.063162
risk_nn	0.118782	0.121072
savage_nn	0.059428	0.128837
stab_nn	0.143965	0.123195
stroke_vb	0.118052	0.089683
thump_nn	0.096097	0.104708
tip_vb	0.092821	0.085674
tree_nn	0.026129	0.035296
twist_nn	0.111025	0.117444
word_nn	0.034585	0.046221

Table 5: JSD (within each target word and time period) between the frame distributions derived from the lemmatised and raw corpora.

## C Predictions of FTFE model

True Positive	True Negative	False Positive	False Negative
prop_nn	tree_nn	quilt_nn	ball_nn
graft_nn	face_nn		head_nn
stab_nn	risk_nn		ounce_nn
plane_nn	bag_nn		
tip_vb			
twist_nn			

Table 6: True Positive, True Negative, False Positive, False Negative results of FTFE.