

EPIC-EuroParl-UdS: Information-Theoretic Perspectives on Translation and Interpreting

Maria Kunilovskaya, Christina Pollkläsener

Saarland University, University of Hildesheim
maria.kunilovskaya@uni-saarland.de, pollklaesener@uni-hildesheim.de

Abstract

This paper introduces an updated and combined version of the bidirectional English–German EPIC-UdS (spoken) and EuroParl-UdS (written) corpora containing original European Parliament speeches as well as their translations and interpretations. The new version corrects metadata and text errors identified through previous use, refines the content, updates linguistic annotations, and adds new layers, including word alignment and word-level surprisal indices. The combined resource is designed to support research using information-theoretic approaches to language variation, particularly studies comparing written and spoken modes, and examining disfluencies in speech, as well as traditional translationese studies, including parallel (source vs. target) and comparable (original vs. translated) analyses. The paper outlines the updates introduced in this release, summarises previous results based on the corpus, and presents a new illustrative study. The study validates the integrity of the rebuilt spoken data and evaluates probabilistic measures derived from base and fine-tuned GPT-2 and machine translation models on the task of filler particles prediction in interpreting.

Keywords: information theory, translation and interpreting, cross-modality, word alignment, parallel corpus, surprisal, LLMs

1. Introduction

In recent years, information-theoretic approaches have gained traction in translation and interpreting studies as a way to explore processing effort (Carl and Schaeffer, 2017; Teich et al., 2020; Wei, 2022) or difficulty (Lim et al., 2024) and linguistic variation, especially triggered by modality (written or spoken) (Nikolaev et al., 2020; Lapshinova-Koltunski et al., 2022; Przybyl et al., 2022a; Kunilovskaya et al., 2025). Yet, researchers in this area often lack ready-to-use resources annotated with word-level information-theoretic indices. Therefore, they either rely on corpus frequencies, i.e., unigram probabilities that ignore context, or must generate such data themselves, which can be time-consuming and resource-intensive.

The updated EPIC-UdS and EuroParl-UdS corpora, combined into EPIC-EuroParl-UdS corpus for English↔German (DE↔EN, i.e., DE-EN, EN-DE), address this gap by providing word-level surprisal values derived from both off-the-shelf (or base) and fine-tuned neural models. The underlying GPT-2 data have been used in several studies, demonstrating the usefulness of surprisal for addressing questions in process-oriented interpreting research and in studies of linguistic variation across spoken and written modes (Kunilovskaya et al., 2025; Pollkläsener et al., 2025; Pollkläsener and Kunilovskaya, 2025). This paper extends the corpus by adding per-word surprisal from machine translation (MT) models and examines their performance on predicting filler particles (FPs, *eah*, *hum*, *hm*) in interpreting.

The combined corpus builds on the previously

released versions Karakanta et al. (2018) for the written subcorpus and Przybyl et al. (2022b) for the spoken subcorpus, recycling their content to varying degrees.

Section 2 places the updated corpus among existing translation and interpreting corpora and introduces key concepts and benefits of information theory. Section 3.1 describes the data structure, formats and availability of the corpus. Section 3.2 explains how the prior work was taken into account and how the current release differs from previous versions. Sections 3.3 and 3.4 provide details on preprocessing and annotation. Section 4 presents applications of the new corpus, including a novel study using and evaluating the new annotation layers (MT surprisal, word alignment, fine-tuned surprisal) to predict FPs in interpreting.

2. Background

Speeches from the European Parliament have long served as a rich data source for translation and interpreting studies. The EuroParl Corpus (Koehn, 2005) remains one of the most influential written parallel resources for numerous language pairs, even though it only provides scarce metadata. EuroParl-UdS (Karakanta et al., 2018) tries to remedy this by re-crawling the speeches and adapting the data for translationese studies, offering a filter based on the speakers' nationality as an approximation for their native language. Amponsah-Kaakyire et al. (2021) use the same re-crawling-pipeline as Karakanta et al. (2018) but align text paragraph-wise, which impedes fine-grained anal-

ysis. Our updated EPIC-EuroParl-UdS corpus includes document-, sentence- and word alignment. On the spoken side, the EPIC corpora family (Bernardini et al., 2018), which includes EPIC-UdS (Przybyl et al., 2022b), relies on manual transcriptions from audio files with European Parliament speeches. Recently, Automatic Speech Recognition (ASR) technologies have been used to accelerate transcription (Fedotova et al., 2024). However, disfluencies transcription still requires manual revision since disfluencies are not yet reliably detected by ASR systems. Our updated corpus is unique in its incorporation of surprisal values for both translation and interpreting data and their sources.

Surprisal is a measure rooted in information theory (Shannon, 1948) that captures the (un)expectedness of a given word based on its probability in context according to a language model. It is considered to be an indicator of cognitive effort required to comprehend or produce language (Hale, 2001; Teich et al., 2020). Unexpected (high surprisal) words are associated with, e.g., longer word durations (e.g. Malisz et al. 2018), longer reading times (e.g. Wilcox et al. 2023) and more disfluencies (e.g. Polkläsener et al. 2025).

In translation and interpreting studies, surprisal and other information-theoretic measures have been used in two types of process studies. First, in experimental settings, surprisal has been correlated with behavioural data such as reading times, keystroke logging, and predictive eye movements (e.g. Carl and Schaeffer, 2017; Hodzik and Williams, 2017; Wei, 2023; Lim et al., 2024). Second, in naturalistic production studies, these measures have been used to annotate and analyse texts in parallel and comparable corpora (e.g., Martínez and Teich 2017; Nikolaev et al. 2020; Lapshinova-Koltunski et al. 2022; Kunilovskaya et al. 2023).

Most of these studies use n-gram surprisal, including unigram surprisal, estimated from the very corpus under analysis. It is our hope that our corpus will enable information-theoretic research using more robust modelling approaches, in which the corpus under analysis serves as held-out test data rather than the source of probability estimates.

3. Corpus Design and Construction

This section describes how previous EuroParl-based resources were updated and extended to create *EPIC-EuroParl-UdS*. This update was motivated by several factors.

First, we aimed to ensure consistent annotation and formatting across the spoken and written components to eliminate any preprocessing-related discrepancies. The previous release of EuroParl-UdS lacked a non-tokenised raw version, and EPIC-

UdS appeared without punctuation or interpreter IDs. The two components had dissimilar file- and segment-naming conventions, data formats and contained incomplete metadata.

Second, we excluded all English–Spanish data, making the new resource strictly bidirectional English↔German, as the Spanish data were not comparable in size or directionality across modalities.

Third, the written and spoken components shared overlapping material (speeches appearing in both forms), with the degree of overlap varying substantially across translation directions: 145 speeches from DE-EN spoken subcorpus also appear in written, while in EN-DE only five documents are shared. To prevent bias in direction-wise cross-modal comparisons, we removed overlapping documents from the larger written component.

Fourth, the EuroParl-UdS data were imbalanced across translation directions, with DE–EN containing roughly 55% more segment pairs. Such an imbalance introduces bias in translationese studies, where sources in one translation direction are used as comparable originals for targets in the other direction. In the new corpus, the written component is balanced across translation directions. Additionally, we improved the automatic sentence alignment in the written subcorpus by using domain-specific glossaries and introducing a quality cut-off value to exclude poorly aligned document pairs. The resulting alignment quality was verified manually.

Finally, we considered extending the written corpus with debates translated after July 2018, the date of the latest speech included in the previous release. However, the new data could not be retrieved using the existing EuroParl-UdS extraction pipeline without compromising the continuity of the original data collection principles. Moreover, as of 1 November 2023, the European Parliament ceased to publicly release official translations as part of the verbatim reports.¹

3.1. Structure and Availability

The EPIC-EuroParl-UdS Corpus comprises two subcorpora—*written and spoken*—each distributed in three complementary data formats.² Each data format is a set of compressed plain-text, UTF-8 encoded tabular files (.tsv.gz), where rows correspond to textual units, and columns contain the associated text, metadata, annotations, or numeric values. Each format emphasises a different anno-

¹https://www.europarl.europa.eu/doceo/document/RULES-9-2023-07-10-RULE-204_EN.html

²Corpus: <https://zenodo.org/records/19010272>; DOI: 10.5281/zenodo.19010272

Code repository: <https://github.com/SFB1102/b7-lrec2026>

tation layer, while item identifiers remain consistent across all tables.

vertical This word-level format is accepted as primary and is based on Stanza tokenisation (see Section 3.3). It includes `word_id`, 10 standard UD annotation fields, word-level surprisal values from GPT-2 and MT models, and word alignments, along with a number of key metadata useful for sorting the data.

long In this format, each row represents a segment (`seg`) and columns store the corresponding metadata, as well as segment-level GPT-2 surprisal aggregates, i.e. average segment surprisal values (`gpt_AvS`).

wide This format provides a parallel, segment-level view for each subcorpus and features aggregates of MT indices (`mt_AvS`).

Some columns may contain N/A values including predominantly string-based fields such as `token`, `pos`, `lemma`, etc. in the vertical format. These arise from alignment asymmetries (e.g. omissions or additions in translation/interpreting) where no corresponding material exists in the source/target. N/A values are also used for words that do not receive surprisal scores, such as FPs, expanded contractions, or surprisal annotation errors.

The item ID pattern is `<ttype>_<mode>_<src_lang>_<tgt_lang>_<doc_id>-<seg_id>_<word_id>`. For example, `ORG_SP_DE_EN_131-02:001` refers to the first word (001) in the second segment (02) of the original (ORG) German document (131), interpreted into English within the spoken subcorpus (SP).

The `speaker_id` attribute warrants special attention. These identifiers are particularly valuable for regression-based analyses, where speaker and/or document IDs can be incorporated as random effects to account for inter-speaker and inter-document variability. In the written subcorpus, the `speaker_id` field corresponds to the `name`; however, the identities of translators are not available (N/A is filled for all absent values). In the spoken subcorpora, the `speaker_id` field contains speaker names for original speeches, whereas interpreters are identified by newly introduced numeric codes (e.g., `fDE1`) based on manual voice analysis and gender-based verification.

The spoken corpora in all formats appear in two versions: a `clean` version without FPs and segment pairs where either source or target was empty and a full version (`fp_none`) including these features.

Additional document-level metadata inherited from earlier releases (e.g. `nationality`, `birth_date`, `birth_place`, `n_party`, etc.) are made available in one separate file per subcorpus.

Appendix A has a snapshot of the EPIC-EuroParl-UdS corpus file structure (Figure 2) as well as lists of text and metadata fields associated with each format (see Table 5).

3.2. Data Collection, Bitext and Filtering

Initial text extraction and bitext creation are described separately for the spoken and written data, as each data source format required a number of subcorpus-specific steps.

Interpreting subcorpus: EPIC-UdS. The current corpus release introduces newly realigned and punctuated data based on the previously manually transcribed material, preserving eight types of disfluencies (see Przybyl et al., 2022b for details). Adding punctuation improved comparability with written data and enabled more reliable downstream processing. A fragment from a spoken transcript with phonetic annotations of disfluencies is shown below:

and finally, / *hum* / I'm [1#I am] seeking to / *euH*
take out / the s/ [s:] the ad/ dition [2#addition] of
split and *hm* separate [s:perate] v/ ow/ votes
[v:otes] [3#] / to [to:] the procedure that will
permit / the President to refer *euH* / back to a
[a:] *euH* / committee, / a r/ f/ f/ f/ f/ report /
which has attracted m/ ow/ m/ more [4#] than
euH f/ fifty [f:ifty] [2#] substantive a/ a/ a/ am/
m/ mendments [6#amendments].

A preliminary analysis of the distribution of annotation categories indicated that some categories were inconsistently transcribed. For the sake of reliability, only FPs were preserved in the combined corpus. However, the segment-level spoken data includes the total counts of all annotated disfluencies (`disfluencies`),³ counts for the three more reliably annotated categories, namely `FPs`, `midword breaks`, `repetitions` and `truncations` (`fillers+3`) and counts of FPs only (`fillers`, including counts of *euH*, *hum*, *hm* forms).

After text preprocessing (standardising forms, mapping phonetic variants and removing all annotated disfluencies except FPs), the segment appears in the released corpus as follows:

And finally, *hum* I am seeking to *euH* take out
the addition of split and *hm* separate votes to
the procedure that will permit the President
to refer *euH* back to a *euH* committee, a re-
port which has attracted more than *euH* fifty
substantive amendments.

Beyond disfluencies, another spoken-specific characteristic concerns segment alignment. Unlike its written counterpart, the spoken data contains empty source or target segments, reflecting

³Monospace script is used throughout the paper to refer to column names in the released tabular data.

cases where interpreters omitted or added content. Depending on the experimental setup, empty or very short segments (fewer than N words) can be treated differently. For experiments where parallelism between source and target is not required (e.g., comparing DE originals and DE targets), such segments can be filtered independently from originals and targets. For experiments requiring strict parallelism, the entire source-target pairs have to be dropped where either side is empty or too short. This will result in a dataset up to 10% smaller than the initial spoken subcorpus. The quantitative parameters of spoken-specific features are summarised in Table 1.

lpair	ttype	segs	% empty	FP	% segs_w/FP
DE-EN	source	3,249	2.2	632	12.6
	target		8.5		2,328
EN-DE	source	3,440	2.4	1,217	23.4
	target		7.5		3,255

Table 1: Ratios of empty segments and segments with filler particles (FPs), and total counts of FPs in spoken data by text type and translation direction.

Segments containing FPs account for more than one-third of interpreted segments in both translation directions, with an average of approximately two FPs per segment.

Translation subcorpus: EuroParl-UdS. The officially released written source speeches and their translations (up to July 2018) were re-extracted from XML files downloaded from the European Parliament website using the EuroParl-UdS scripts.⁴ We also applied the existing filter to retain only original speeches delivered by native speakers of English or German, ensuring greater reliability of translationese comparisons between original and translated language.

The extracted data were re-aligned using LF Aligner (Farkas, 2015), a wrapper around the *Hunalign* library (Varga et al., 2008), with domain-specific bilingual glossaries built from *IATE* dictionaries (European Union, 2024).

The resulting parallel corpus was restricted to documents with an average *document-level* alignment score above 0.3 for DE-EN and 0.5 for EN-DE.

In automatically aligned written data, empty segments typically reflect alignment errors rather than genuine translation features. Documents containing aligned empty source (*src_seg*) or target (*tgt_seg*) segments were excluded, except when such segments occurred at the beginning or end of a document or contained three words or fewer; in those cases, only the segment was removed. This filtering affected 1,235 document pairs in DE-EN

direction and 1,072 document pairs in EN-DE (i.e., about 10% in each direction), most of which were removed entirely.

A manual evaluation of 80 randomly selected document pairs, where originals had approximately the same number of segments (750 segment pairs per direction), conducted by a compensated research assistant, showed misalignment rates below 4.5% for DE-EN and 1.8% for EN-DE.

Given the strong manual evaluation results and the much larger size of the written subcorpus (compared to spoken), testing alternative alignment algorithms such as *Vecalign* (Thompson and Koehn, 2019), which mainly benefits low-resource cases, was deemed unnecessary.

3.3. Text Standardisation

The same set of text preprocessing operations was applied uniformly to all texts before any word-level annotation. These operations were guided by our prior experience with the data and by the intended use of the corpus for cross-modal, cross-lingual, and monolingual comparisons, enabling consistent analysis across both translation directions.

The preprocessing steps included the removal of non-printable ASCII characters (except tabs and newlines), standardisation of single and double quotes, dashes and superscripts. All hyphenated words were force-tokenised to improve comparability across languages (e.g., *EVP-Fraktion*, *well-known*).

3.4. Annotation

With this corpus release, we provide three layers of word-level annotation: standard morphological and syntactic annotation, word surprisals derived from base and fine-tuned GPT-2 and MT models, and word alignments.

Stanza parse. Linguistic annotation (tokenisation, POS tagging, lemmatisation, and dependency parsing) was performed using Stanza v1.10.1 (Qi et al., 2020) with default pre-trained models for English and German, corresponding to the Universal Dependencies treebanks distributed with this release. Segment alignment was preserved throughout the parsing process. Multi-sentence segments (about twice as frequent in the spoken subcorpus) were parsed sentence by sentence and retained within the same aligned translation unit by maintaining continuous word numbering in *word_ids* across each segment. The main difference between annotating the spoken and written subcorpora was that, for spoken data, FPs had to be removed prior to annotation and then reinserted at their original positions into the parsed segment with FP tag using the *word_ids* index. The vertical

⁴<https://github.com/hut-b7/europarl-uds>

CoNLL-U output of *Stanza* was used as the canonical tokenisation scheme used in the subsequent indexing with LLM-based surprisals and for word alignment, subject to an important caveat concerning *multitokens* discussed below.

Multitokens are multiword tokens, most notably English contractions (*it’s*, *there’s*, *I’d*) and Genitive case forms and syncretic German prepositions such as *am* and *zur*. As illustrated in Appendix A (Figure 6), these forms are represented simultaneously at two levels in the data. On one level, surface multiword tokens are preserved for surprisal annotation and word alignment, which operate on the actually realised forms. On the other level, the internal morphosyntactic structure of these contractions is made explicit through multitoken expansion (e.g., *it’s*→*it is*, *im*→*in dem*, *December’s*→*December’s*), which is required for fully functional syntactic dependency trees. The true surprisal values exist only for surface realisations; the surprisal values in the rows containing expansions are set to N/A, while the surface token `word_id` receives an additional identifier (e.g. `SI_DE_EN_030-21:010:1` for *there in there’s*).

To enable comparisons across modes and translation directions, the written subcorpus was divided into *train* and *test* splits. Prior to this split, any speeches occurring in both written and spoken modes were removed from the written data.

The *test* split consists of documents selected to match the spoken subcorpus in size and document length (at least 12 segments per document; 170 document pairs per translation direction), ensuring comparable segment counts. The spoken corpus was assigned entirely to the *test* split.

The *train* split was balanced across translation directions by randomly subsampling the larger German–English portion to match the English–German data in terms of segment count. This split was used to fine-tune the LLMs for surprisal estimation. The written data are released in separate tables corresponding to the train and test splits. The resulting data split parameters are presented in Table 2.

To further characterise the data, we report additional descriptive parameters. Average segment length ranges from 24.1 to 26.8 words for written data and 18.6 to 20.7 words for spoken data, with standard deviations of approximately 14 words. Minimum segment length is not controlled for and can be as short as one word. Maximum segment length reaches 317 words in the written German target, where multiple sentences are merged into a single translation segment.

In spoken data, targets contain roughly twice as many multi-sentence segments as sources (DE–EN: 8.4% vs. 17.4%; EN–DE: 7.8% vs. 19.9%). In written data, multi-sentence segments are rare overall (1.1–7.5%), and in con-

mode	lpair	type	words	segs	docs
spoken	DE-EN	src, org	63,172	3,249	165
		tgt	61,826		
	EN-DE	src, org	70,130	3,440	137
		tgt	63,727		
written (test)	DE-EN	src, org	76,327	3,218	170
		tgt	85,053		
	EN-DE	src, org	75,804	3,070	170
		tgt	75,668		
written (train)	DE-EN	src, org	2,924,739	122,407	10,367
		tgt	3,292,542		
	EN-DE	src, org	3,110,670	122,233	11,690
		tgt	3,102,631		

Table 2: Corpus statistics for *EuroParl-EPIC-UdS*, based on vertical format. Spoken data counts include FPs and empty segments on either side. Multiword tokens are counted once as surface-level items. Abbreviations: lpair=language pair, src=source, org=original, tgt=target; docs=documents, segs=segments.

trast to spoken, shows a clear directional asymmetry: English-to-German favours sentence splitting, while German-to-English favours sentence merging, implying denser information packaging in written English. For the remainder of the paper, we treat the segment—rather than the sentence—as the operational unit, although the vast majority of aligned segments contain a single sentence on either the source or the target side.

Surprisal indexing. Surprisal is a key measure of information in the information-theoretical approach to language study. It is calculated as an inverse base-2 log probability of a word in context (Equation 1):

$$S(w) = -\log_2(P(w|context)) \quad (1)$$

The probabilistic measures in this study are derived from monolingual GPT-2 small models and neural MT models. The GPT-2 small family was selected because its surprisal estimates have been shown to correlate best with experimental measures of processing difficulty (Oh and Schuler, 2023). Each word (defined as a *Stanza* token in vertical corpus format) in the EPIC-EuroParl-UdS test subsets is annotated with surprisal values obtained from language-specific pre-trained (base) and fine-tuned GPT-2 models (English Radford et al. (2019); German Schweter (2020)). In addition, each target word is assigned surprisal values derived from pre-trained and fine-tuned, translation-direction-specific MT models, developed by Tiedemann and Thottingal (2020). The vocabulary size of all four models is approximately 50 K tokens. Fine-tuning was conducted exclusively on the written train split of the corpus (comprising written orig-

inals and written parallel data for the MT baseline). Consequently, the spoken targets represent an out-of-domain (OOD) challenge for the MT models. For the GPT-2 models, this challenge is twofold: the targets are OOD in terms of both mode (spoken vs. written) and text type (original vs. targets). Detailed training hyperparameters are documented in Table 7 (Appendix C), followed by a comprehensive evaluation of the resulting fine-tuned models. The training process yielded expected results: the cross-entropy loss for fine-tuned models is lower than for base models even on out-of-domain test sets (except `deen_mt`, see Figure 5); the training and validation losses go down with no signs of overfitting (Figure 6); subcorpus-averaged surprisal and BLEU is better for fine-tuned than for base models (Table 8) and individual examples demonstrate that domain specific vocabulary is less surprising for fine-tuned models (Figure 7). Train subsets are included in the corpus and are indexed with surprisal from base models only.

Word-level surprisal is computed within segment boundaries. For each subword, the probability is conditioned on the left context up to the segment boundary, explicitly marked with `<|bos|>` token. Rare, longer segments in the training dataset are truncated to 150 leftmost subwords. Consequently, surprisal is estimated with a varying-length left context, capped at 150 tokens. This setup was compared with a 64-subword sliding-window approach. The two methods were strongly correlated, yielding linear regression R^2 between 0.74 (spoken German targets) and 0.81 (written English targets) (see Figure 4, Appendix B). Because the sliding-window approach produced comparable patterns, albeit on a reduced scale (as illustrated in Figure 3, Appendix B), it was not pursued further, and the segment-bounded model was adopted for all subsequent analyses.

Our approach to input formatting and post-processing was guided by two principles: maintaining the naturalistic quality of the source text and ensuring a robust alignment between subword-level surprisals and linguistic tokens.

Unlike standard pipelines that rely on expanded syntactic representations, we opted for naturalistic raw text input. This involved detokenising punctuation, preserving hyphenated tokens as single units, and utilising surface multiword tokens (e.g., `zur`, `he'll`, `Council's`) rather than their expanded Universal Dependencies (UD) forms (e.g., `zu der`, `he will`, `Council 's`). This was critical for accuracy: preliminary tests showed that feeding fully tokenised segments inflated probabilities for punctuation and multiword tokens, as the underlying language models were pre-trained on raw text.

LLM subword probabilities were converted to log space and pre-aggregated based on the model's

beginning-of-word markers (e.g., GPT-2's \hat{G}), while keeping punctuation as separate units. However, this \hat{G} -based pre-aggregation did not guarantee an exact match with Stanza tokens. In approximately 3% of written test segments, mismatches occurred where a single pre-aggregated unit spanned multiple Stanza tokens (e.g., Stanza: `['über', '99', '%', '.']`; GPT-2: `['¼ber', '99', '%.']`) or vice versa (Stanza: `['at', '5.30', 'p.m.', 'on', 'Thursday']`; GPT-2: `['at', '5.30', 'p.m.', '.', 'on', 'Thursday']`). These cases appeared almost exclusively around complex punctuation sequences. Rather than discarding these segments and compromising the dataset, we implemented a sequential set of rules to recover values for common edge cases, including abbreviations (i.e., `p.m.`), float-like numbers (20 000, 0,7%-Zielgröße), and punctuation sequences (`'%,'`, `'IV.,'`, `':'`). When subword values required splitting to align with Stanza tokens, $0.75 \times$ value was assigned to the preceding token and $0.25 \times$ value to the trailing punctuation mark; alternatively, subword surprisals were summed.

Despite our efforts to recover most erroneous segments, a few dozen cases in the 120K-segment train splits remained misaligned. Rather than excluding these segments, we indexed them with N/A values. This preserves the original segment alignment and document structure, facilitating future document-level analyses and the annotation of higher-level discourse phenomena, such as coreference, without sequence gaps.

The vertical format of the spoken subcorpus retains FPs and empty segments to preserve the characteristics of spoken language. However, these elements were removed when computing surprisal values. They were reinserted during post-processing at their original positions, with all corresponding surprisal values set to N/A.

Segment-level corpus formats contain average segment surprisal (AvS), computed both at the Stanza token level (excluding FPs) and at the subword level (e.g., for GPT-2 columns: `base_gpt_AvS`, `base_gpt_AvS_subw`, `ft_gpt_AvS`, `ft_gpt_AvS_subw`). Additionally, we included segment-level *pseudo-BLEU* scores (`base_bleu`, `ft_bleu`) calculated between the reference target segment and token-wise argmax predictions under gold target prefixes, using *sacrebleu* (Post, 2018). To avoid overly harsh zero scores when higher-order n-grams do not match, we apply exponential smoothing (`smooth_method="exp"`). This approach allows sentence-level BLEU to reflect partial literal matches, making it more stable and informative than strict n-gram matching. Higher BLEU scores indicate that the MT model finds the translation more straightforward or typical, suggesting lower cross-lingual transfer difficulty.

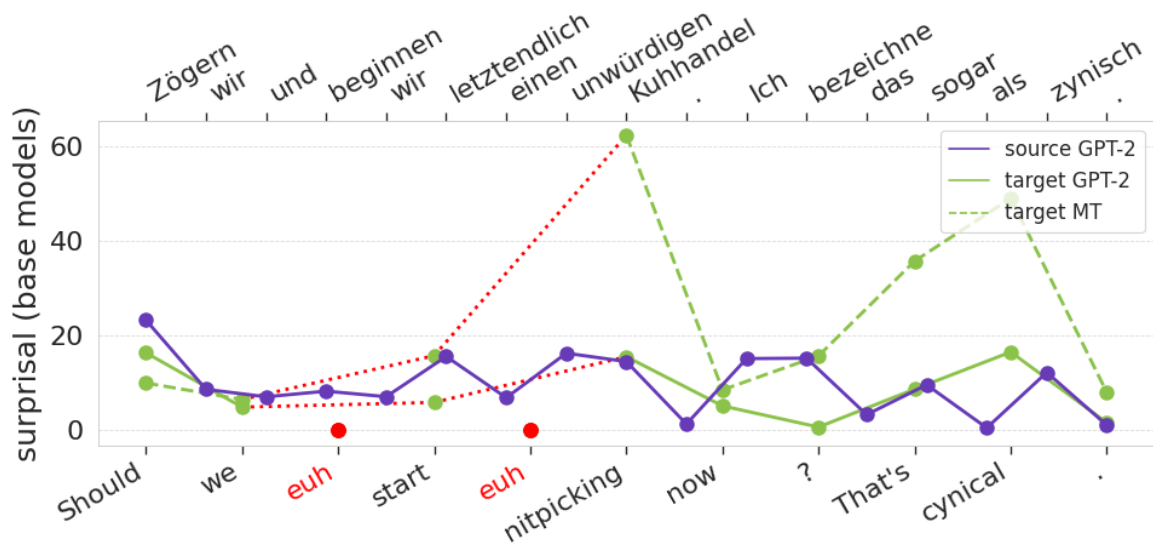


Figure 1: Word surprisals from base monolingual GPT-2 models and base MT model for a segment pair.

Figure 1 shows a spoken parallel segment containing FPs and two sentences (seg_id: `_DE_EN_023-08`). It reports three types of word-level surprisal values: base monolingual GPT-2 surprisals for source and target words, and base cross-lingual MT surprisals for target words.

Table 8 (Appendix B) offers subcorpus-averaged surprisals for each text type and model (both base and fine-tuned). While these averages show only small differences between categories, word-level surprisal captures local phenomena (e.g. FPs), justifying the release of word-level surprisals in this corpus.

A key finding is the divergent behaviour of the monolingual and MT models after fine-tuning. For GPT-2, fine-tuning consistently reduces average surprisal across all subcorpora. Conversely, MT models exhibit higher surprisal after fine-tuning, despite a concurrent increase in pseudo-BLEU scores (indicating better translation quality). Furthermore, mean surprisals are systematically higher for the spoken mode compared to the written mode, likely reflecting the higher entropy and structural irregularities inherent in spoken discourse.

The relationship between GPT-2 and MT indices is considered to capture the trade-off between accuracy and fluency objectives in translation (Teich et al., 2020). Theoretically, translating requires balancing fidelity to the source (low MT surprisal) with fluency in the target (low GPT-2 surprisal). When this balance is hard to achieve, translators can opt for a literal translation strategy (low MT surprisal) at the expense of target language fluency (high GPT-2 surprisal) or deviate from the source (high MT surprisal) in order to comply with target norms (low GPT-2 surprisal). Based on this reasoning,

(Lim et al., 2024) expected a negative correlation between MT and GPT-2 surprisal. Their results, computed over the combined data from all 13 translation directions using multilingual GPT-2 and MT surprisals, confirm the hypothesis. They report a single weak but statistically significant negative correlation between MT and GPT-2 surprisal.

Our data, however, reveals a consistent nonlinear relation between GPT-2 and MT surprisals across subcorpora. The hypothesised negative correlation holds only for simpler segments (below 11 bits/word). Above this threshold, increased deviation from the source do not contribute to target fluency, quite the opposite. Generalised Additive Models (GAM) confirm that these nonlinearities are systematic and predictive (R^2 0.17–0.28), suggesting that the hypothesised accuracy-fluency trade-off collapses as cross-lingual transfer difficulty increases (see Figure 8 in Appendix C).

The differences in findings likely stem both from the neural models used for surprisal estimation and from the nature of the data. Unlike prior work based on small-scale, experimentally elicited student translations, our study draws on professionally produced, high-quality translations in a comparatively more demanding domain, which may fundamentally alter the observed accuracy–fluency dynamics. We further argue that the degree of nonlinearity in cross-lingual transfer captured by MT surprisal reflects the scale of production effort required by a given segment. Once this effort exceeds a certain threshold, neither translators nor interpreters can consistently maintain audience-oriented fluency, pointing to inherent constraints of the production process itself.

mode	lpair	src_tokens	% unaligned	% multi-aligned
spoken	DE-EN	60,125	34.96	4.42
spoken	EN-DE	66,729	34.48	2.14
written	DE-EN	76,327	19.59	5.08
written	EN-DE	75,804	20.03	1.49

Table 3: Word alignment statistics for test sets source tokens. Unaligned and multi-aligned columns have ratios of source tokens not aligned or aligned to several target tokens.

Word alignment. Inspired by AWESoME (Dou and Neubig, 2021), we use the *bert-base-multilingual-cased* model (Devlin et al., 2019) to extract contextualised subword representations. For every pair of source and target subword embeddings, we compute dot-product similarity and apply bidirectional softmax normalisation (source→target and target→source). Subword pairs whose mutual bidirectional softmax score exceeds the recommended threshold of 0.01 are marked as aligned; this value constitutes the alignment confidence threshold used throughout the paper.

Aligned subword pairs are then aggregated to the word level using the pre-computed subword–word maps. For each word pair, the final alignment score is obtained as the mean of the bidirectional softmax scores, averaged across all contributing subword pairs. If a source word exceeds the same alignment confidence threshold at the word level, its corresponding target word(s) are recorded in the columns `aligned_word` and `aligned_word_id`.

The softmax-based alignment permits one-to-many mappings; in such cases, these columns contain lists of target words and their IDs. As shown in Table 3, this word alignment pattern is direction-dependent, occurring in over 4% of source words in the DE–EN data and in about 2% in EN–DE, irrespective of mode. In contrast, the proportion of unaligned source words (marked with an explicit ‘NA’ string in the data tables) reflects mode-specific variation rather than language-pair differences. As expected, the spoken subcorpora exhibit a higher proportion of unaligned source words (up to 35%), consistent with the observation that interpreting tends to yield less literal, less isomorphic renditions.

4. Applications of Information Theory

This section presents possible applications of the corpus. While section 4.1 outlines a novel study, section 4.2 refers to a related study and describes ideas for further possible applications.

4.1. Interpreting: FP Prediction

In this study, we look at the link between FPs and surprisal. Previous work has shown that FPs are followed by high-surprisal words, (Zámečník, 2019; Dammalapati et al., 2021; Raof et al., 2024), indicating upcoming lexical retrieval difficulty. Our previous study (Pollkläsener et al., 2025) replicated this effect in interpreting. However, we focused solely on the effect of target language surprisal (formulation difficulty). We now extend this work using the new annotation to include source surprisal (comprehension difficulty), MT surprisal (transfer difficulty), word alignment, and a comparison of base and fine-tuned models.

Using mixed-effects logistic regression, we model if a target word is preceded by a FP (1) or not (0). FPs make up roughly 5% of all target words in both EN–DE and DE–EN (3,255 and 2,328 instances, respectively). Predictors represent different types of processing difficulty: source surprisal (comprehension), target surprisal (formulation), and MT surprisal (transfer). Both local (word-level) and global (segment-average) measures of the three surprisal types are included in the model as global surprisal contributes explanatory value (Pollkläsener and Kunilovskaya, 2025). Every target word in the corpus is annotated with target and MT surprisal. To obtain comprehension surprisal, we use word- and segment-alignment. Multi-aligned target words are averaged and unaligned target words are excluded. Separate models are estimated for EN–DE and DE–EN directions as English and German surprisal values are not directly comparable.

First, we compare fine-tuned and base surprisal. To do so, we fit models using either base or fine-tuned surprisal values for each direction. In both directions, models with base surprisal have lower AIC values than fine-tuned surprisal models (DE–EN: base AIC = 11359, fine-tuned AIC = 11473; EN–DE: base AIC = 15313, fine-tuned AIC = 15349), indicating better model fit. Similarly, models with base surprisal achieve higher C-scores, reflecting better discrimination between outcomes (DE–EN: base C = 0.6994, fine-tuned C = 0.686; EN–DE: base C = 0.6515, fine-tuned C = 0.6432). We conclude that base surprisal values perform better in this task.

Table 4 presents the fixed-effect estimates for the models using base surprisal values. Given the consistency of the results across directions, we present a joint discussion below. Formulation surprisal measures (`nxtwS_tgt`, `AvS_tgt`) emerge as the strongest predictors for FP occurrence, which is noteworthy, as we had expected transfer difficulty to be the primary indicator of production difficulty in interpreting. One possible explanation is that MT surprisal is computed using more source context

	DE-EN		EN-DE	
	Coef	SE	Coef	SE
Intercept	-3.49*	0.18	-3.04*	0.08
nxtwS_tgt	0.54*	0.03	0.37*	0.03
nxtwS_src	-0.10*	0.03	-0.14*	0.03
nxtwS_mt	0.04	0.03	0.08*	0.02
AvS_tgt	-0.35*	0.04	-0.40*	0.03
AvS_src	0.17*	0.03	0.20*	0.03
AvS_mt	-0.20*	0.03	-0.03	0.03

Table 4: Fixed-effect estimates (Coef) and standard errors (SE) from mixed-effects logistic regression models predicting FPs in the target. Surprisal values are from base models. Predictors are z-scored. Asterisk indicates statistical significance ($p < .05$).

(the entire aligned source segment) than the interpreter may have heard at a given moment. Future studies could take into account time alignment and estimate surprisal based only on the source context available to the interpreter up to the last word heard.

Interestingly, the coefficient for average target surprisal is negative, indicating that lower average segment surprisal is associated with higher FP probability. This contrasts with the findings reported in [Pollkläsener and Kunilovskaya \(2025\)](#), though considering all three surprisal measures together may clarify this pattern. For next-word surprisal, formulation and transfer difficulty (nxtwS_tgt, nxtwS_mt) show positive effects, whereas comprehension difficulty (nxtwS_src) has a negative effect. Interpreters thus tend to produce FPs before words that are difficult to formulate and transfer but easy to comprehend. We see the opposite for the global measures. Average target and MT surprisal (AvS_tgt, AvS_mt) show negative effects, while average source surprisal (AvS_src) has a positive effect on FP occurrence. This pattern suggests that comprehension difficulties may not surface at the level of the immediately produced word. Instead, they appear to manifest more globally, with processing demands accumulating over the segment. Moreover, interpreters appear to compensate for the high cognitive load of retrieving the next word by producing material with low transfer and formulation difficulty in the rest of the segment.

4.2. Cross-Modal Studies

In [Kunilovskaya \(2026\)](#), EPIC-EuroParl-UdS is used to examine translationese. The study asks whether translation task difficulty can explain the degree to which translations deviate from comparable original texts in the target language, assuming that producing translationese- or interpretese-ridden text requires less effort. Translationese is quantified as the segment-level probability of being

a translation (*translatedness* score) returned by a binary classifier. This score as the response variable in a regression experiment. Translation task difficulty is separated into comprehension difficulty, operationalised by a set of source-related predictors such as source text surprisal or lexical density, and transfer difficulty, operationalised by predictors reflecting the relationship between source and target, such as MT surprisal or mean alignment score. The results show that transfer difficulty is more effective in predicting translatedness than comprehension difficulty (except in spoken DE-EN). Information-theoretic predictors (e.g. surprisal, entropy) match or outperform traditional features (e.g. tree depth, lexical density) in written mode, but offer no advantage in spoken mode.

EPIC-EuroParl-UdS can also be used to examine the influence of V-final clauses in German sources on English spoken and written targets. While translators can access the entire source text before producing their output, interpreters render segments in real time, often having to begin their interpretation before the sentence has fully unfolded. [Hodzik and Williams \(2017\)](#) investigated the prediction of sentence-final verbs in simultaneous interpreting from German into English. They found that more predictable verbs (as indicated by contextual semantic cues) were associated with shorter latency times, suggesting facilitated production through predictability. This effect could be investigated using EPIC-EuroParl-UdS, relating the surprisal of the clause-final verbs in the German source to the surprisal of the aligned target words to reveal transfer strategies in translation and interpreting.

5. Conclusion

This paper presents the surprisal-annotated, updated English↔German translation and interpreting corpus *EPIC-EuroParl-UdS* and its possible applications. We correct and refine the corpus content and structure, and introduce new annotation layers designed for information-theory-motivated studies, especially cross-lingual ones, at the local level. The corpus pre-processing pipeline emphasises document integrity by resolving issues with individual segments instead of skipping them. By providing this resource, we hope to facilitate advances in contrastive and translationese studies, especially since the vertical format is directly openable in *R*, simplifying direct data analysis. All data and code are available under a creative commons licence. We plan to further improve *EPIC-EuroParl-UdS* by incorporating surprisal measures from other models (e.g., multilingual GPT-2 and LLaMA) and adding time-alignment.

6. Acknowledgements

This research was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1102 Information Density and Linguistic Encoding, Project-ID 232722074.

7. Bibliographical References

- Kwabena Amponsah-Kaakyire, Daria Pylypenko, Cristina España-Bonet, and Josef van Genabith. 2021. [Do not rely on relay translations: Multilingual parallel direct Europarl](#). In *Proceedings for the First Workshop on Modelling Translation: Translatology in the Digital Age*, pages 1–7, online. Association for Computational Linguistics.
- Silvia Bernardini, Adriano Ferraresi, Mariachiara Russo, Camille Collard, and Bart Defrancq. 2018. *Building Interpreting and Intermodal Corpora: A How-to for a Formidable Task*, pages 21–42. Springer Singapore, Singapore.
- Michael Carl and Moritz Jonas Schaeffer. 2017. [Why translation is difficult: A corpus-based study of non-literality in post-editing and from-scratch translation](#). *HERMES - Journal of Language and Communication in Business (Denmark)*, (56):43–57.
- Samvit Dammalapati, Rajakrishnan Rajkumar, Sidharth Ranjan, and Sumeet Agarwal. 2021. [Effects of duration, locality, and surprisal in speech disfluency prediction in English spontaneous speech](#). In *Proceedings of the Society for Computation in Linguistics 2021*, pages 91–101.
- 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, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou and Graham Neubig. 2021. [Word alignment by fine-tuning embeddings on parallel corpora](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2112–2128, Online. Association for Computational Linguistics.
- Alice Fedotova, Adriano Ferraresi, Maja Miličević Petrović, and Alberto Barrón-Cedeño. 2024. [Constructing a multimodal, multilingual translation and interpreting corpus: A modular pipeline and an evaluation of ASR for verbatim transcription](#). In *Proceedings of the 10th Italian Conference on Computational Linguistics (CLiC-it 2024)*, pages 349–355, Pisa, Italy. CEUR Workshop Proceedings.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Proceedings of the Second Meeting of the North American Chapter of the Association for Computational Linguistics on Language Technologies*, page 1–8, USA. Association for Computational Linguistics.
- Ena Hodzik and John N Williams. 2017. [Predictive processes during simultaneous interpreting from German into English](#). *Interpreting*, 19(1):1–20.
- Alina Karakanta, Mihaela Vela, and Elke Teich. 2018. [EuroParl-UdS: Preserving and extending metadata in parliamentary debates](#). In *Proceedings of the LREC 2018 Workshop ParlaCLARIN: Creating and Using Parliamentary Corpora*, pages 33–38, Miyazaki, Japan.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *Proceedings of Machine Translation Summit X: Papers*, pages 79–86, Phuket, Thailand.
- Maria Kunilovskaya. 2026. [Translationese as a Rational Response to Translation Task Difficulty](#). Preprint.
- Maria Kunilovskaya, Sharid Loáiciga, and Ekaterina Lapshinova-Koltunski. 2025. [Information divergence in translation and interpreting: Findings from same-source texts](#). In *KONVENS: 21th Conference on Natural Language Processing (KONVENS 2025) ; Proceedings of the Conference. Volume 1: Long and Short Papers*, pages 192–202. HsH Applied Academics.
- Maria Kunilovskaya, Heike Przybyl, Elke Teich, and Ekaterina Lapshinova-Koltunski. 2023. [Simultaneous interpreting as a noisy channel: How much information gets through](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*, pages 608–618, Varna, Bulgaria. INCOMA Ltd.
- Ekaterina Lapshinova-Koltunski, Christina Pollkäsener, and Heike Przybyl. 2022. [Exploring explicitation and implicitation in parallel interpreting and translation corpora](#). *The Prague Bulletin of Mathematical Linguistics*, 119:5–22.
- Zheng Wei Lim, Ekaterina Vylomova, Charles Kemp, and Trevor Cohn. 2024. [Predicting human translation difficulty with neural machine](#)

- translation. *Transactions of the Association for Computational Linguistics*, 12:1479–1496.
- Zofia Malisz, Erika Brandt, Bernd Möbius, Yoon Mi Oh, and Bistra Andreeva. 2018. [Dimensions of segmental variability: Interaction of prosody and surprisal in six languages](#). *Frontiers in Communication*, 3.
- José Manuel Martínez Martínez and Elke Teich. 2017. [Modeling routine in translation with entropy and surprisal: A comparison of learner and professional translations](#). In Larissa Cercel, Marco Agnetta, and Maria Teresa Amido Lozano, editors, *Kreativität und Hermeneutik in der Translation*, pages 403–427. Narr Francke Attempto Verlag.
- Dmitry Nikolaev, Taelin Karidi, Neta Kenneth, Veronika Mitnik, Lilja Saeboe, and Omri Abend. 2020. [Morphosyntactic predictability of translationese](#). *Linguistics Vanguard*, 6(1):1–12.
- Byung-Doh Oh and William Schuler. 2023. [Why does surprisal from larger transformer-based language models provide a poorer fit to human reading times?](#) *Transactions of the Association for Computational Linguistics*, 11:336–350.
- Christina Polkläsener and Maria Kunilovskaya. 2025. [Euh... where do interpreters hesitate? An information-theoretic perspective on sentence-initial filler particles in simultaneous interpreting](#). In *12th edition of the Disfluency in Spontaneous Speech Workshop (DiSS 2025)*, pages 92–96.
- Christina Polkläsener, Maria Kunilovskaya, and Elke Teich. 2025. [Surprisal explains the occurrence of filler particles in simultaneous interpreting](#). *SKASE Journal of Translation and Interpretation*, 18(2):55–78.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Heike Przybyl, Alina Karakanta, Katrin Menzel, and Elke Teich. 2022a. [Exploring linguistic variation in mediated discourse: Translation vs. interpreting](#). In Marta Kajzer-Wietrzny, Adriano Ferraresi, Ilmari Ivaska, and Silvia Bernardini, editors, *Mediated discourse at the European Parliament: Empirical investigations*, chapter 7, pages 191–218. Language Science Press, Berlin.
- Heike Przybyl, Ekaterina Lapshinova-Koltunski, Katrin Menzel, Stefan Fischer, and Elke Teich. 2022b. [EPIC-UdS - creation and applications of a simultaneous interpreting corpus](#). In *Proceedings of the 13th Conference on Language Resources and Evaluation (LREC 2022)*, pages 1193–1200, Marseille, France. ELDA.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. 2020. [Stanza: a Python natural language processing toolkit for many human languages](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). *OpenAI*.
- Negin Raoof, Yating Wu, Carlos Bonilla, Junyi Jessy Li, Stephanie M Grasso, Alex Dimakis, and Zoi Gkalitsiou. 2024. [Modeling bilingual disfluencies with large language models](#). In *ICML 2024 Workshop on LLMs and Cognition*.
- Claude E. Shannon. 1948. [A mathematical theory of communication](#). *Bell System Technical Journal*, 27:379–423.
- Elke Teich, José Martínez Martínez, and Alina Karakanta. 2020. [Translation, information theory and cognition](#). In Fabio Alves and Arnt Lykke Jakobsen, editors, *The Routledge Handbook of Translation and Cognition*, 1st edition, pages 360–375. Routledge.
- Brian Thompson and Philipp Koehn. 2019. [Vecalign: Improved sentence alignment in linear time and space](#). 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 1342–1348.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – building open translation services for the world](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Dániel Varga, Péter Halácsy, András Kornai, Viktor Nagy, László Németh, and Viktor Trón. 2008. [Parallel corpora for medium density languages](#). In *Recent advances in natural language processing IV: selected papers from RANLP 2005*, pages 247–258. John Benjamins Publishing Company.
- Yuxiang Wei. 2022. [Entropy as a measurement of cognitive load in translation](#). In *Proceedings of the 15th Biennial Conference of the Association for Machine Translation in the Americas. Empirical Translation Process Research*, pages 75–86.

- Yuxiang Wei. 2023. *Ambiguity and Entropy in the Process of Translation and Post-editing*. PhD thesis, Dublin City University.
- Ethan G Wilcox, Tiago Pimentel, Clara Meister, Ryan Cotterell, and Roger P Levy. 2023. [Testing the predictions of surprisal theory in 11 languages](#). *Transactions of the Association for Computational Linguistics*, 11:1451–1470.
- Jiří Zámečník. 2019. *Disfluency prediction in natural spoken language*. PhD thesis, Dissertation, Universität Freiburg, 2019.
- Tiedemann, Jörg and Thottingal, Santhosh. 2020a. *OPUS-MT English-to-German Translation Model*. Hugging Face Model Hub. PID <https://huggingface.co/Helsinki-NLP/opus-mt-en-de>.
- Tiedemann, Jörg and Thottingal, Santhosh. 2020b. *OPUS-MT German-to-English Translation Model*. Hugging Face Model Hub. PID <https://huggingface.co/Helsinki-NLP/opus-mt-de-en>.

8. Language Resource References

- Jacob Devlin and Ming-Wei Chang and Kenton Lee and Kristina Toutanova. 2018. *BERT multilingual base model (cased)*. Hugging Face Model Hub. PID <https://huggingface.co/google-bert/bert-base-multilingual-cased>.
- European Union. 2024. *IATE: Interactive Terminology for Europe*. European Union Publications Office. PID <https://iate.europa.eu/search/standard>.
- Farkas, András. 2015. *LF Aligner tool*. <https://sourceforge.net/projects/aligner/>. Version 3.12.
- Karakanta, Alina and Vela, Mihaela and Teich, Elke. 2018. *EuroParl-UdS Corpus*. CLARIN-D-Zentrum repository hosted at Saarland University, Germany. PID <http://hdl.handle.net/21.11119/0000-0000-D5EE-4>.
- Przybyl, Heike and Karakanta, Alina and Menzel, Katrin and Teich, Elke. 2022. *EPIC-UdS Corpus*. CLARIN-D-Zentrum repository hosted at Saarland University, Germany. PID <http://hdl.handle.net/21.11119/0000-0008-F519-8>.
- Qi, Peng and Zhang, Yuhao and Zhang, Yuhui and Bolton, Jason and Manning, Christopher D. 2020. *Stanza: A Python Natural Language Processing Toolkit for Many Human Languages*. Stanford NLP Group. PID <https://stanfordnlp.github.io/stanza/>.
- Radford, Alec and Wu, Jeffrey and Child, Rewon and Luan, David and Amodei, Dario and Sutskever, Ilya. 2019. *GPT-2 English Model*. Hugging Face Model Hub. PID <https://huggingface.co/openai-community/gpt2>.
- Schweter, Stefan. 2020. *German GPT-2 Model*. Hugging Face Model Hub. PID <https://huggingface.co/dbmdz/german-gpt2>.

Appendix A. Details on Corpus Formats

--spoken	--written
--epic-clean_long.tsv.gz	--europ-test_long.tsv.gz
--epic-fp-nones_long.tsv.gz	--europ-train_ende_long.tsv.gz
--epic-clean_wide.tsv.gz	--europ-train_deen_long.tsv.gz
--epic-fp-nones_wide.tsv.gz	--europ-test_wide.tsv.gz
--epic-clean_vrt.tsv.gz	--europ-train_ende_wide.tsv.gz
--epic-fp-nones_vrt.tsv.gz	--europ-train_deen_wide.tsv.gz
--epic_org_doc_meta.tsv.gz	--europ-test_vrt.tsv.gz
	--europ-train_ende_vrt.tsv.gz
	--europ-train_deen_vrt.tsv.gz
	--europ_org_doc_meta.tsv.gz

Figure 2: Structural overview of the EPIC-EuroParl-UdS repository, illustrating the organisation of spoken and written sub-corpora across data formats (long, wide, and vertical) and associated metadata.

Format	Field groups	Fields
Vertical (word level)	Modelling factors CoNLL-U Surprisal indices Word alignment Context	doc_id, seg_id, word_id, lpair, lang, mode, ttype, speaker_id id, token, lemma, pos, xpos, rel, feats, head_id, deps, misc srp_base_gpt2, srp_base_mt, srp_ft_gpt2, srp_ft_mt (except training sets) aligned_word_id, aligned_word raw_seg (except training sets)
Long (segment level)	Modelling factors GPT-2 surprisal indices Responses (spoken) Context	doc_id, seg_id, lpair, lang, mode, ttype, speaker_id, <i>spoken</i> : delivery_rate, delivery_wpm, speech_timing_sec, source_text_delivery_type base_gpt_AvS, base_gpt_AvS_subw, ft_gpt_AvS, ft_gpt_AvS_subw (except training sets) disfluencies, fillers, fillers+3 raw_seg, tokens, wc_tok
Wide (segment-pair)	Core identifiers Text fields MT surprisal and BLEU	src_doc_id, src_seg_id, lpair, mode src_raw_seg, tgt_raw_seg base_mt_AvS_subw, base_mt_AvS, base_bleu, ft_mt_AvS_subw, ft_mt_AvS, ft_bleu (except training sets)

Table 5: Data fields available in each release format. All formats are provided as gzipped UTF-8 tab-separated tables (.tsv.gz).

word_id	id	head_id	token	srp_base_gpt2	srp_ft_gpt2	srp_base_mt	srp_ft_mt	aligned_word	speaker_id
..._030-21:001	<NA>	<NA>	It's	13	12.5	35.3	38.4	Begonnen, ist	fEN3
..._030-21:001:1	1	5	It	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3
..._030-21:001:2	2	5	's	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3
..._030-21:002	3	5	all	5.9	6.4	15.8	19.7	<NA>	fEN3
..._030-21:003	<NA>	<NA>	euh	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3
..._030-21:004	<NA>	<NA>	hm	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3
..._030-21:005	<NA>	<NA>	euh	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3
..._030-21:006	4	5	very	6.5	1.6	7.6	7.5	guten	fEN3
..._030-21:007	5	0	well-intended	16.7	19.5	45.8	46.4	guten, Absichten	fEN3
..._030-21:008	6	5	.	3	3.4	10.2	9	<NA>	fEN3
..._030-21:009	1	3	But	2.2	1.3	9.9	15.9	leider	fEN3
..._030-21:010	<NA>	<NA>	there's	5.6	10.7	43	52.1	<NA>	fEN3
..._030-21:010:1	2	3	there	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3
..._030-21:010:2	3	0	's	<NA>	<NA>	<NA>	<NA>	<NA>	fEN3

Table 6: Excerpt from the vertical-format data illustrating the handling of multiword tokens (e.g. *it's*, *there's*), where surface forms used for alignment and surprisal annotation coexist with expanded morphosyntactic representations required for syntactic parsing (seg_id: SI_DE_EN_030-21). Surprisal values are rounded to 0.1. Other UD and metadata fields are omitted.

Appendix B. Sliding Window Results




Figure 3: Word surprisals from base GPT-2 using segment-bounded and sliding window approaches (seg_id: SI_EN_DE_129-27)

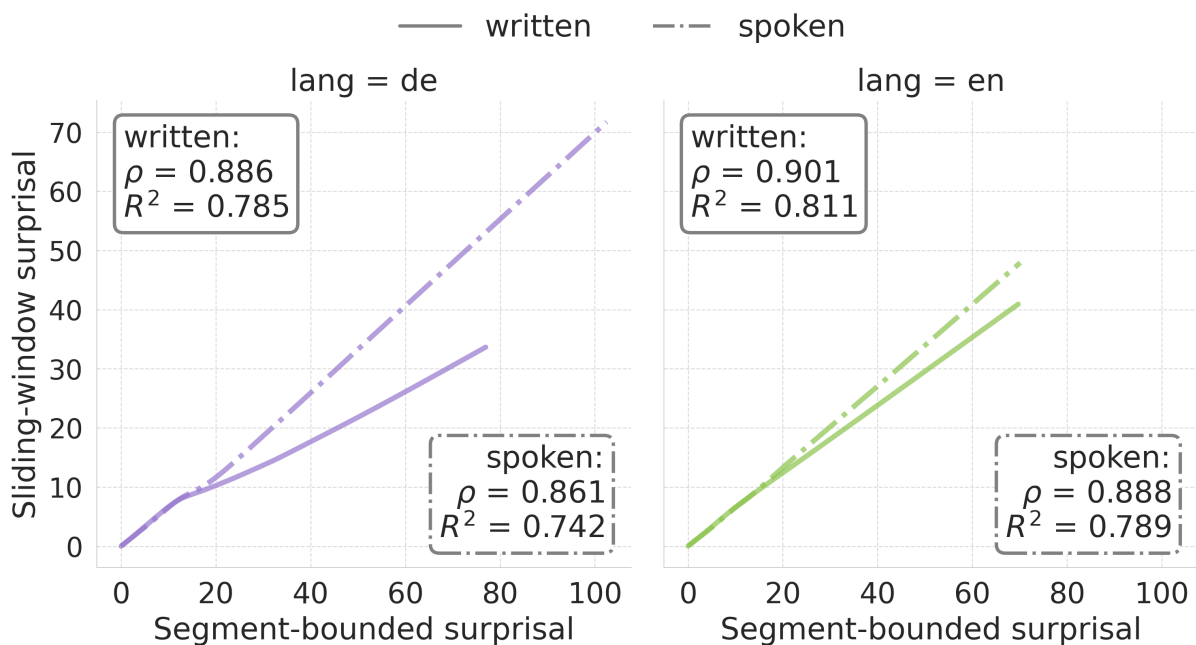


Figure 4: Relation between segment-bounded and sliding window base GPT-2 word surprisals for written and spoken target subcorpora by language. Lines represent LOWESS-smoothed fits (frac = 0.2).

Appendix C. Fine-Tuning Details

Hyperparameter	GPT-2	MT
batch_size	32	16
gradient_accumulation_steps	2	3
warmup_ratio	0.05	0.10
epochs	10	10
learning_rate	8×10^{-6}	8×10^{-6}
weight_decay	0.01	0.001
early_stopping δ	0.0005	0.001
early_stopping_patience	3	3
max_sequence_length	150	150
label_smoothing	0.0	0.1
deterministic_training	False	False

Table 7: Training hyperparameters for GPT-2 and MT systems.

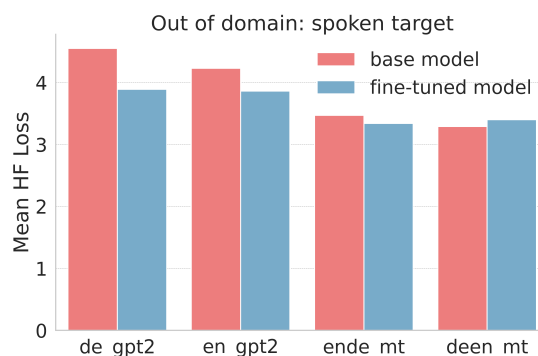
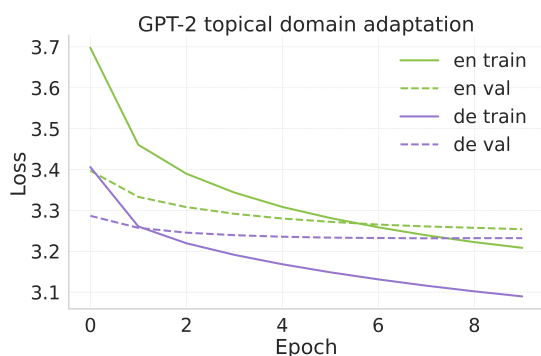



Figure 5: Hugging Face cross-entropy loss on out-of-domain test sets.



(a) GPT-2 losses



(b) MT losses

Figure 6: Training and validation loss curves for finetuning experiments. (a) GPT-2 base model, (b) MT baseline. The curves have a slow adaptation profile due to a conservative learning rate (8×10^{-6}) and weight decay. While validation progress is gradual, this parameterisation was necessary to prevent overfitting.

			mean UD token surprisal				subcorpus BLEU	
mode	lpair	ttype	base_gpt2	ft_gpt2	base_mt	ft_mt	base_bleu	ft_bleu
spoken	deen	source	6.997	5.929	–	–	–	–
		target	6.419	5.863	12.195	13.641	12.508	12.826
	ende	source	6.488	5.804	–	–	–	–
		target	7.137	6.102	12.480	13.819	11.064	12.450
written	deen	source	6.222	5.185	–	–	–	–
		target	5.949	4.951	11.099	12.348	27.146	29.550
	ende	source	6.051	4.899	–	–	–	–
		target	5.991	5.203	12.728	14.116	23.038	28.177

Table 8: Base vs. fine-tuned surprisal and pseudo-BLEU scores for EuroParl-EPIC-UdS test sets. UD-token surprisal and BLEU scores improve after fine-tuning (as expected), but cross-lingual surprisal from MT models is worse.

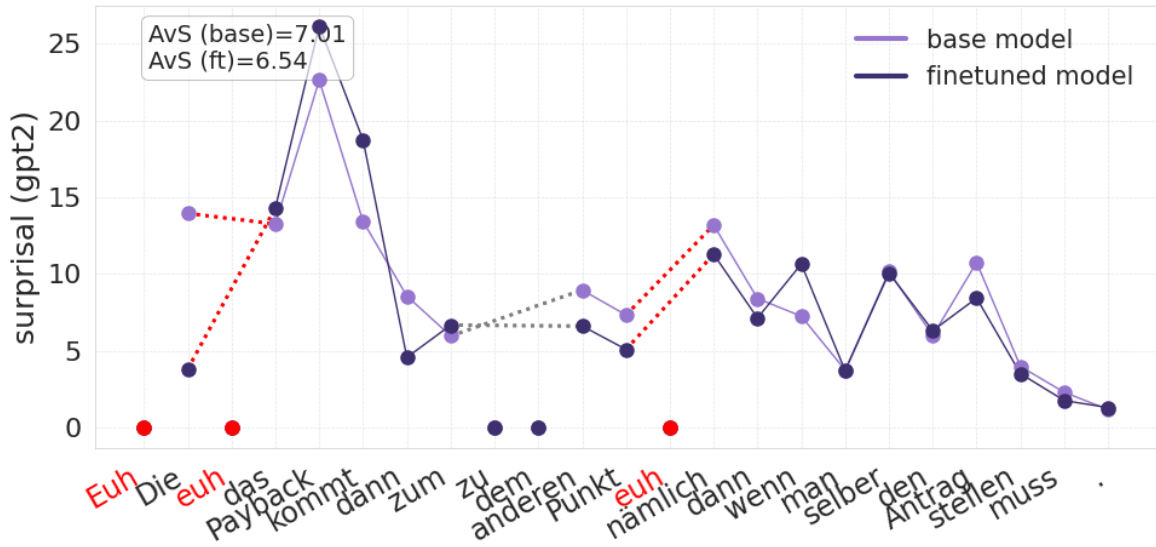


Figure 7: Word-surprisals from base and fine-tuned GPT-2 models for a target with filler particles and expanded UD multiwords (seg_id: SI_EN_DE_129-27). Domain-specific vocabulary (e.g. *Antrag*) is less surprising for fine-tuned models, while general language words (e.g. *kommt*) are more surprising compared to base models.

Appendix D. Extended Case Study Illustrations

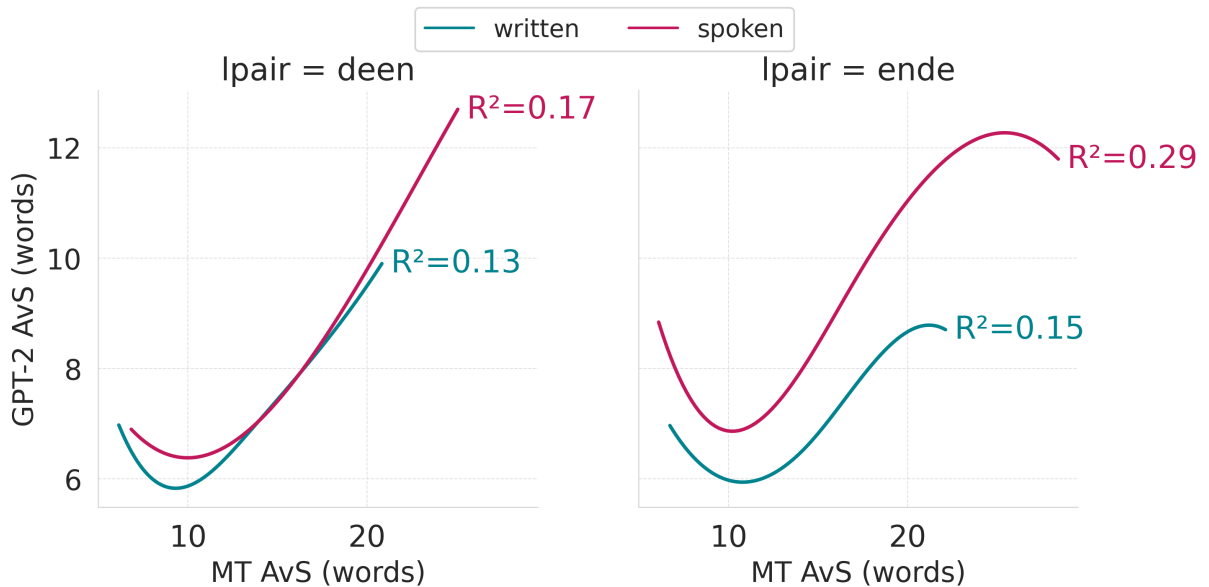


Figure 8: Fitted linear GAM regression lines and pseudo R^2 statistics. Setup: $n_splines=5$, gridsearch run on word surprisals from base models for test sets in EuroParl-EPIC-UdS.