

Thesis proposal: Are We Losing Textual Diversity to Natural Language Processing?

Josef Jon and Ondřej Bojar

Charles University, Faculty of Mathematics and Physics

Institute of Formal and Applied Linguistics

Prague, Czech Republic

jon@ufal.mff.cuni.cz

Abstract

This thesis argues that the currently widely used Natural Language Processing algorithms possibly have various limitations related to the properties of the texts they handle and produce. With the wide adoption of these tools in rapid progress, we must ask what these limitations are and what are the possible implications of integrating such tools into our daily lives.

As a testbed, we have chosen the task of Neural Machine Translation (NMT). Nevertheless, we aim for general insights and outcomes, applicable to current Large Language Models (LLMs). We ask whether the algorithms used in NMT have inherent inductive biases that are beneficial for most types of inputs but might harm the processing of untypical texts, thereby contributing to a cycle of monotonous, repetitive language – whether generated by machines or humans. To explore this hypothesis, we define a set of measures to quantify text diversity based on its statistical properties, like uniformity or rhythmicity of word-level surprisal, on multiple scales (sentence, discourse, language). We conduct a series of experiments to investigate whether NMT systems struggle with maintaining the diversity of such texts, potentially reducing the richness of the generated language, compared to human translators.

We further analyze potential origins of these limitations within existing training objectives and decoding strategies. Ultimately, our goal is to propose and validate alternative approaches (e.g., loss functions, decoding algorithms) that maintain the diversity and complexity of language and that allow for better global planning of the output generation, enabling the models to better reflect the ambiguities inherent in human communication.

1 Introduction

Language technologies powered by machine learning, such as predictive typing, machine translation, text generation, or chat assistants have become

deeply integrated into our daily lives. With the advancement of Large Language Models (LLMs), we can expect that interaction with such tools will become an even more integral part of our work and social interactions. Yet, many important questions regarding these tools remain unanswered.

Do they understand language the same way we do? Can they capture the full richness, creativity, and diversity inherent in human communication? Are there any biases within the algorithms themselves that can be beneficial for processing ordinary types of texts, but harmful for specific cases that deviate from the usual rules found in mundane texts?

As our reliance on these technologies grows, we risk adapting and simplifying our language to fit their capabilities. Could this lead to further proliferation of simplistic, mundane text, which will be in turn used to train a new generation of models, making them even less adept at processing surprising inputs? Will machine-produced, monotonous content dominate our informational space, causing authentic texts to become lost within it?

Ultimately, could this trend contribute to a loss of the diversity and richness of human language and, consequently, human thought?¹

This thesis explores such questions through the lens of one specific application of these language technologies: Neural Machine Translation (NMT). We investigate NMT's performance on texts that diverge from the norm not in terms of terminology or domain but in the structure and organization of their information content. Our aim is to identify intrinsic properties of texts that current NLP systems consistently struggle with, regardless of their scale.

Even if we answered all the proposed questions positively, identifying failure cases and adverse impacts of current tools, it would not change the rate at which these technologies are adopted. Thus, we propose innovative decoding algorithms and train-

¹For example, [Kranich \(2014\)](#) explains how language contact results in language changes

ing strategies to preserve linguistic diversity and ensure that language technology enhances, rather than limits, our expressive potential.

2 Problems and proposed solutions

We briefly summarize the goals of the thesis in this section. The literature describing the characteristics we discuss here is listed in Section 3.

2.1 Measuring diversity

First, we establish criteria to assess the uniformity and typicality of the distribution of information on multiple scales, including the sentence level, the discourse level, and the whole structure of the language itself. Our methodology is based on the existing literature (described in Section 3), and our objective is to associate the measures at each level with observable real-world phenomena.

At the word level, our exploration into measuring information content – or surprisal – examines its association with word-level reading times. We use surprisal estimates obtained by language models to predict the reading times. For sentences and paragraphs, we assess the uniformity of these surprisal values. We evaluate different methodologies based on their correlation and their predictive capacity for sentence-level reading times and linguistic acceptability as judged by humans.

At the discourse level, we intend to observe rhythmic patterns in surprisal distribution to gauge engagement by predicting at what point the audience may stop reading an article or listening to a podcast. We hypothesize that a periodical change between quick and slow pace of information transmission has an effect on the reader’s enjoyment.

Finally, in the context of entire languages, we will estimate the optimal information rate or channel capacity language-wide, including comparisons across different languages.

The properties we intend to observe during the translation process at different levels of the language are also listed in Figure 1.

Developing the measures for these properties will allow us to:

- 1) Identify potentially problematic or difficult types of texts for NMT,
- 2) evaluate the current algorithms and our innovations,
- 3) identify texts that were already translated by MT in the wild.

Points 1) and 2) are integral to this thesis, and point 3) helps to counter one of the challenges that

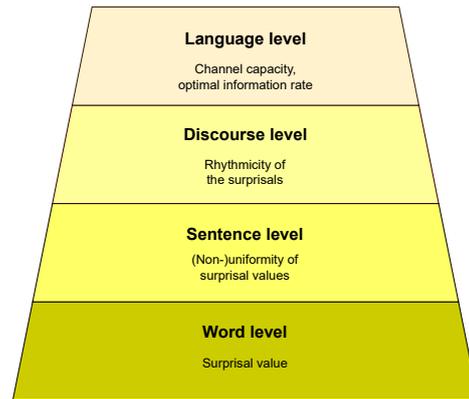


Figure 1: Proposed text properties that we plan to observe during the MT process.

is already present in MT and arises more with the continuing adoption of LLMs: new models training on the outputs of previous generations of models due to pollution of the Web by machine-generated text. Such a self-feedback loop inevitably leads to loss of diversity over multiple iterations, up to a completely degenerated output (Guo et al., 2023; Briesch et al., 2023).

Section 4.1 discusses our progress in this task.

2.2 Diversity in MT

We employ the uniformity measures explored in the previous part to answer one of the initial questions: *Does NMT make the surprisal distribution in the translation more uniform than a human translator would?* We gather a diverse set of datasets to look for correlations in surprisal distribution between source and both human and machine translation. We also compare the absolute values of the measures between professional human translation and MT. As we discuss in Section 4.2, the results we have gathered so far are not conclusive, as many factors play a role in computing the surprisal. However, we did find some evidence of MT failing to maintain non-uniformity of surprisal distribution in highly surprising texts, like poetry.

2.3 Causes of excessive uniformity

Assuming we will, in fact, find that the NMT struggles to produce non-uniform text, we will investigate the causes of this behavior. One of the principal suspects is the beam search decoding. The beam search was introduced as an approximation of the maximum-a-posteriori (MAP) decoding, which is in practice intractable. However, it has been shown, quite surprisingly, that it produces better results than the exact MAP decoding in the context of NMT (Stahlberg and Byrne, 2019), suggesting that

it introduces some own inductive bias into the translation process. Meister et al. (2020) show that this bias is linked to uniformity – beam search prefers hypotheses where the surprisal is uniformly distributed and the same results can be obtained with exact MAP decoding with uniformity regularizer.

We hypothesize that this regularization is necessary to produce high-quality MT outputs because of insufficiencies in the modeling. The model is predominantly trained to generate the next word given the source text and the previously produced target text, a process that does not inherently encourage global planning strategies. Without beam search enforcing uniformity of the surprisals, the model would make surprising decisions that it could in theory "balance out" in the future, but, given the lack of global planning, it performs this long-range balancing poorly. Beam search masks this inadequacy of the model by enforcing surprisal uniformity.

This leads us to believe that to replace beam search in order to improve the translation of non-uniform texts, we should also look into training objectives and improve the global planning capabilities of the model.

2.4 Alternative decoding algorithm

We propose an alternative decoding algorithm for NMT. Sampling-based algorithms are commonly used in Large Language Models (LLMs), but they have been shown unfit for traditional NMT with smaller model sizes. However, Minimum Bayes risk decoding (MBR) with sampled hypotheses has been very successful recently and we extend this approach by using a genetic algorithm to combine and modify the translation candidates, see Section 4.3. We are circumventing the problem of too uniform surprisals by using MT quality evaluation and estimation metrics to guide the decoding, ignoring the probability estimated by the model.

In the future, we will also address decoding in LLMs and compare the properties of LLM-generated text to the conventional NMT models.

2.5 Alternative training objective

In line with the previous reasoning, we have also investigated alternative training objective functions for NMT to allow for better global planning of the translation. The intuitive reasoning behind this is that global planning allows the model to estimate how much probability mass the best solution has from the start, so it can plan for more uneven distribution of it throughout the output sequence, taking

more "risks" and selecting higher surprisal tokens, since it knows it will get the probability "back" in the future timesteps. We assume that beam search mitigates this behavior, which is beneficial for models where the global planning contains many errors. Better global planning could thus, among other benefits, remove the need for beam search "fixing" the produced translation.

The approaches can range from simply predicting more future tokens at a single timestep (Pal et al., 2023; Stern et al., 2018), through sentence-level objectives (Goyal et al., 2019; Lu et al., 2020), to modeling hierarchical structures of the text (Ainslie et al., 2020).

The long-term planning abilities are also inherently related to the intrinsic uncertainty of the translation task, stemming from the fact that each source sentence has multiple correct translations. Current modeling approaches are inadequate since they only model a single probability distribution over the target sentence (although viable alternatives exist; Stahlberg and Kumar (2022)). In left-to-right decoding, this uncertainty rises with the distance of the future we are trying to predict, as at each next token, the word choices are split into many possible, and all correct, pathways. Thus, the inadequacy of the modeling becomes even more pressing in our distant future modeling use case.

We propose using an auxiliary training objective, Contrastive Predictive Coding (CPC; van den Oord et al. (2019)), that forces the internal representations of the model at each timestep to be more similar to future internal representations on multiple levels of hierarchy (words, phrases, sentences, and paragraphs). This allows us to model the future, while simultaneously alleviating the uncertainty issue, since instead of forcing a prediction of a single token, we predict an internal representation of a more abstract nature at some point in the future. However, in current models, the internal representations are still only a single-point estimates of the corresponding words or other linguistic phenomena. We will explore approaches to model the uncertainty in embeddings themselves, similar to Kesiraju et al. (2020).

Conversely, the past context is also modeled inadequately in NMT. Traditional *teacher forcing* practice, which feeds the previous reference token into the decoder during training, imposes a single "correct" past context on the model. This causes *exposure bias*: the model never sees its own gen-

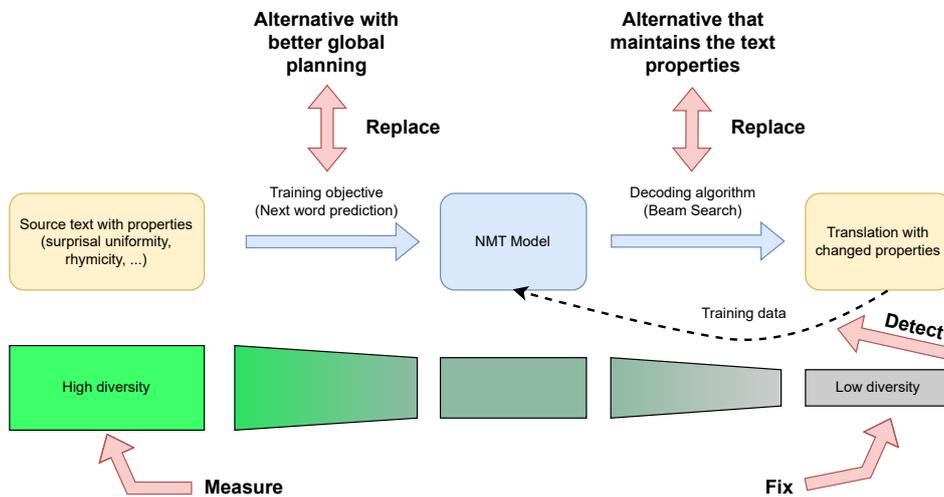


Figure 2: Illustration of the current NMT process and the points addressed in this thesis (red arrows).

erated prefixes during the training, but they are used at test time (Ranzato et al., 2016). This bias negatively affects the fidelity of the statistical properties of generated text to the training data. In a way, this discourages the model from planning into the future (or accounting for possible futures) when generating the current token, since in the next step, the current token is thrown away anyway. To address this, we will look into alternatives that, for instance, introduce a lattice of possible synonyms rather than a single reference token, offering a range of potential past contexts, or we will alternate between model’s predictions and gold tokens during the training. Additionally, we aim to explore sentence-level training objectives that permit the model to generate complete hypotheses freely, evaluating them semantically against the reference, inspiring ourselves in modern MT quality estimation approaches.

Reinforcement learning from human feedback (RLHF; Stiennon et al. (2022); Ouyang et al. (2022)) has recently gained popularity with LLMs and it can be also viewed as a sentence-level objective. Later in our work, we will explore it along with other techniques used in LLMs.

See Section 4.4 for details of the work on training objectives carried out so far.

2.6 Objectives

Based on the previous observations, we set out to reach 5 objectives:

- **O1:** Create a simple measure to score the non-uniformity of the surprisal distribution on multiple levels in order to identify texts that are inherently challenging for today’s NMT systems.
- **O2:** Identify the root causes of the problems: How do different steps of the NMT process affect the properties of the output?

- **O3:** Develop alternative decoding algorithms that mitigate the issues.
- **O4:** Develop alternative training objectives that allow for better global planning.
- **O5:** Perform a case study where the resulting system produces better translations in a real-world scenario.

We illustrate the NMT process, the potential problems we posit, and the places where we intervene to counter them (red arrows) in Figure 2.

3 Characteristics of NMT-produced language

In this section, we summarize the related work on properties of translations produced by NMT and the origins of these properties.

3.1 Lexical diversity in NMT

Lexical diversity, reflecting the richness of vocabulary, is commonly assessed using metrics like Type-Token Ratio (TTR) and Measure of Textual Lexical Diversity (MTLD; McCarthy, 2005). Studies highlight that Neural Machine Translation (NMT) systems often produce less lexically diverse outputs compared to human translations, frequently simplifying language or reinforcing biases.

Research by Vanmassenhove et al. (2019, 2021) demonstrates that MT systems generally favor frequent words and neglect rare vocabulary. Similarly, Toral (2019) observed reduced lexical diversity in both raw and post-edited machine translations. In contrast, Junczys-Dowmunt (2020) noted comparable lexical diversity between NMT and human translations in news articles, although such texts typically require less creativity. Brglez and Vintar (2022) initially reported higher lexical diversity in NMT outputs but clarified this was often due to

errors, emphasizing that assessments of diversity must always account for translation accuracy.

3.2 Skewed word frequencies

Lexical uniformity in NMT arises largely because outputs tend to favor high-frequency words, reducing the occurrence of rarer vocabulary (Ott et al., 2018; Koehn and Knowles, 2017). Although sampling-based decoding could potentially mirror the training data distribution more accurately, this approach often compromises translation quality. Combining sampling with Minimum Bayes Risk (MBR) decoding and advanced evaluation metrics can improve results, yet MBR itself can inadvertently reinforce biases toward frequent words due to their greater prevalence in sampled translations (Müller and Sennrich, 2021).

3.3 Decoding Algorithms

Decoding algorithms are the mechanism through which a subsequent word $y(t)$ is predicted in an output sequence at time step t , based on a probability distribution over the target vocabulary, conditioned upon the source sentence x and the sequence of previously generated target words, formalized as $P(y(t)|x, y(t-1), \dots, y(0))$. The probability of a whole target sentence is factorized by the chain rule as a product of these word-level probabilities. The Maximum A Posteriori (MAP) decoding is traditionally employed to identify the most probable translation under the model’s distribution. However, the search space for the decoding is very large: the size of the target vocabulary raised to the power of the maximum sentence length (theoretically unbounded). Therefore, an approximation of MAP is used, most commonly the beam search (Och and Ney, 2004; Graves, 2012). Both beam search and MAP in general have many shortcomings which are discussed in recent literature.

Stahlberg and Byrne (2019) introduce an exact decoding algorithm utilizing depth-first search that can be executed within reasonable time constraints for most sentences. Despite achieving optimal model probabilities, a significant portion of the translations produced end up as empty strings or low quality translations. This discrepancy suggests two conclusions. First, there might be an issue with how the whole translation task is modeled, since often the global optimum of these probabilistic models is an empty string. It also suggests beam search’s effectiveness may not be due to its accuracy as a MAP approximation, but rather due to its

ability to overcome flaws in the modeling approach, challenging the reliance on the model’s probability as the selection criterion.

Meister et al. (2020) also explore the reasons why beam search produces higher quality outputs than an exact MAP search. They propose that the beam search itself introduces an inductive bias. To find it, they reverse engineer the objective that beam search is a solution for, i.e. they try to build MAP-based algorithm that produces the same results as beam search. Their conclusion is that the exact MAP with a uniformity regularizer which enforces Uniform Information Distribution (UID) (Aylett and Turk, 2004; Fenk and Fenk-Oczlon, 1980; Levy and Jaeger, 2006; Bell et al., 2003; Genzel and Charniak, 2002) behaves the same as beam search. The UID hypothesis posits a preference among language users for utterances that distribute information evenly. In other words, beam search does not only look for the most probable solution but also prefers solutions where the probability is distributed evenly across the whole sentence. This property effectively conceals mistakes in NMT modeling and allows for the production of usable translations, even being dubbed the *beam search blessing* by Meister et al. (2020). Wei et al. (2021) employ a similar regularizer in the training of the model, which led to improved translation quality.

Beam search’s preference for selecting hypotheses that adhere to UID principles at the sentence level suggests a potential conflict for certain text types, particularly those where an element of surprise or non-uniform information distribution is desired. This line of thinking forms the basis for our examination of alternative decoding algorithms.

Work on pathologies of NMT, like Koehn and Knowles (2017) and Stahlberg and Byrne (2019), sparked a debate on whether the standard training objective for machine translation, word-level maximum likelihood estimation, is suitable for modeling the task, and if the issues arise from the definition of the objective itself.

Eikema and Aziz (2020) suggest that the objective function is correct and the issue lies within the MAP decoding. They argue that the mode of the model’s distribution is not an adequate decision rule for a problem of as high dimension as NMT — only a very small probability mass is given to the most probable translation. Their analysis shows that sampling-based approaches produce translations more reflective of the training data’s statistical

properties, although they do not achieve optimal translation quality. They advocate for minimum Bayes risk (MBR; [Goel and Byrne \(2000\)](#)) decoding over beam search, as it covers a broader probability range and better preserves the training data’s statistical properties.

MBR does not aim to identify translations with the highest model probability. Instead, it seeks translations that maximize a chosen utility function, often an MT evaluation metric such as COMET. In theory, the utility function is evaluated against the entire set of all possible translations of the input.

Again, the complete computation is intractable in practice, so sampling-based approximations are used, such as sampling a pool of hypotheses from the model and using them as both candidate translations and references, computing the utility function of each hypothesis with respect to all the others. The same authors discuss strategies to construct the pool of hypotheses in [Eikema and Aziz \(2022\)](#).

MBR started to gain more popularity recently because of advances in MT metrics which can be used as the utility function ([Amrhein and Sennrich, 2022](#); [Freitag et al., 2022](#); [Fernandes et al., 2022](#); [Jon et al., 2022](#); [Jon and Bojar, 2023](#)).

3.4 Training objective

Our discussion so far has already touched upon the appropriateness of the prevailing objective in NMT.

Conventionally, this objective models the translation process as computing the probability of a target translation y given a source sentence x , represented by a single probability distribution $P(y|x)$. In practice, this is implemented by a softmax layer on top of the decoder, computing probability distribution over target vocabulary in each decoding timestep. Such a model implicitly assumes the existence of a single “correct” translation for every source sentence – an assumption that contrasts with the linguistic reality where an extensive amount of valid translations can exist for a single source text. This formulation of the objective forces these valid translations to compete for representation within a single probability distribution. This does not allow the model to distinguish between two types of uncertainty: extrinsic uncertainty, caused by noisy training data, and intrinsic uncertainty, caused by the existence of multiple valid translations of a single source sentence. [Stahlberg et al. \(2022\)](#) identify the intrinsic uncertainty as the main culprit behind multiple pathologies in NMT, including the already

discussed issues with MAP decoding.

SCONES (Single-label Contrastive Objective for Non-Exclusive Sequences; [Stahlberg and Kumar \(2022\)](#)) aims to remove this single correct translation assumption by modeling the translation probabilities separately for each (source sentence, possible translation) pair, so that multiple valid translations from training data can be considered correct at the same time. The results suggest that using SCONES improves translation quality over many language pairs and it alleviates the problems that arise with MAP decoding described earlier – the inadequacy of mode and shifting of the text statistics compared to training data. The related work that we base our approach on is described in 2.5.

4 Progress so far

In this section, we describe the work that we have carried out so far towards the defined objectives.

4.1 O1: Information distribution in language

We approach **O1** by developing a metric capable of identifying texts that pose challenges to current NLP algorithms. We aim to anchor this metric in observable real-world phenomena, ensuring it remains interpretable and meaningful rather than relying on opaque measures, e.g., some “difficulty” scores predicted by other machine learning models. Our goal is to concretely define and understand the characteristics that make certain texts particularly difficult for automated systems to process.

We started by focusing on the surprisal theory [Hale \(2001\)](#), which relates cognitive effort to the surprisal value of words, suggesting that the effort to comprehend a word increases with its unpredictability (surprisal) given the context and on the related Uniform Information Density (UID) theory, proposed by [Levy and Jaeger \(2006\)](#). The UID theory posits that for longer sequences, e.g. sentences, the relationship between surprisal levels and cognitive effort is super-linear, suggesting sentences with evenly distributed surprisal are easier to comprehend. We have replicated (with small deviations in the methodology) the experiments from [Meister et al. \(2021\)](#) that test out the validity of the super-linear relationship by its predictive power for reading time and linguistic acceptability ratings. In practical terms, we computed log-level word perplexities by an LM, applied multiple formulas on the resulting array (including the super-linear formula) and computed the predictive power of the

result. We expected that the sentences with a more uneven distribution of the word-level perplexities will be harder to comprehend.

We have found only partial evidence for this relationship, and other newer works, e.g., (Shain et al., 2024), show that the topic is complex. We discuss our experiments in detail in Appendix A.

4.2 O1 and O2: MT and uniformity of surprisal

We use the uniformity measures introduced in the previous section to confirm one of our initial hypotheses: *Is NMT producing translations that are more uniform in terms of surprisal distribution than a human?* This hypothesis posits that source sentences characterized by highly uneven surprisal distributions would exhibit more uniformity upon translation by MT systems, a phenomenon not expected to occur with human translations. We selected multiple datasets from different domains and measured the uniformity of the word-level perplexities in the source, in the human translation (HT), and in the machine translation (MT).

The initial results have shown that MT, unexpectedly, produced less uniform translations than humans. We hypothesized that this might be due to translation errors in MT, which can result in wrong and thus surprising translations. We tried to filter out low-quality translations using COMET score, but found a bias toward low and uniform surprisal texts in this metric itself. The details of the experiments and the results are described in Appendix B. Overall, we found some evidence for our hypothesis in highly unorthodox texts, like poetry, but improvements to our methodology are necessary.

4.3 O3: Alternative decoding algorithms

Our proposed alternative to the beam search is based on sampling, MBR decoding, external resources like dictionaries, and a genetic algorithm (GA). This method involves applying common GA operations – mutation and crossover – to a set of translation hypotheses, guided by a fitness function based on one or more MT metrics. The overview of the method is presented in Appendix C.

4.4 O4: Alternative training objectives

For Objective 4, we aimed to enhance NMT’s ability for long-range planning within translations. We incorporated the Contrastive Predictive Coding (CPC) objective (van den Oord et al., 2019). This algorithm introduces a regularization component

to the training loss, designed to align internal representations in a selected layer in the current timestep to the representations in future timesteps. One of the strengths of CPC is its focus not on predicting the exact future word, but rather on aligning future representations of words. This partially mitigates the issues with the ambiguity of the future, since embeddings tend to be similar for synonyms that would otherwise be penalized by the strict, single-correct-token prediction.

Traditional loss functions like cross-entropy struggle with predicting such high-dimensional targets like the internal representations directly. Instead, we focus on maintaining as much mutual information as possible between these representations. The concrete implementation of the loss function is discussed in Appendix D.

We have carried out initial experiments with this implementation on MT. Overall, we have not seen improvements in automated MT quality scores yet, but we have only scratched the surface of potential experimental settings.

4.5 Conclusions so far

We have confirmed that surprisal estimates from language models are predictive of reading times on word level. On the sentence level, the discussion on the nature of the relationship between word surprisal distribution and reading times and acceptability ratings is more complex. We have evaluated multiple surprisal distribution uniformity measures and our results show a small preference towards more uniformly distributed surprisals in human acceptability assessment. We used these measures to test our hypothesis that MT produces more uniform translations than a human. Again, the results are mixed and vary greatly with chosen experimental settings (language model, tokenization, normalization, dataset). At best, we can see some evidence for our hypothesis in diverse text types, like books and poetry, but more research is necessary.

To replace the potentially problematic beam search, we have developed a novel decoding algorithm, based on sampling, Minimum Bayes Risk decoding, dictionaries, and a genetic algorithm. We confirmed its effectiveness in 1) improving translation quality and 2) creating adversarial examples for arbitrary MT evaluation metrics.

We have implemented an auxiliary training objective, CPC, which aligns internal representations in each step with representations of future steps, as

well as representations of more abstract text units.

5 Future plans

We have taken the first steps to tackle objectives **O1** through **O4**, defined in the Introduction.

5.1 **O1 and O2: Identification of problematic texts and the root causes**

So far, we have only analyzed the surprisal-related properties of the text on the levels of words and sentences. Our results are not conclusive. We plan to run experiments on more diverse datasets, mostly from the literary domain. We also hypothesize that the fact that the LMs are trained on natural text and we use surprisal estimates from them to evaluate uniformity on both natural and MT-produced text might affect the outcomes. We will train a new LM on a 50:50 mix of both types of texts.

We will evaluate alternative approaches to obtaining the word level (pseudo-)surprisal estimates, in place of language models. One possibility is to create a dataset of texts with binary sentence-level labels saying whether people consider the sentence surprising and train a classification model and use attribution methods (Javorský et al., 2023) to calculate the influence of single words in the decision.

We will move to longer textual units. We assume that on the level of a whole discourse, for example, a book or an article, some rhythmic changes between a quick and a slow pace of transmitting information, play a role in reader engagement and enjoyment of the text. We are in the process of obtaining/creating datasets from podcasts with meta-data about where the user stopped listening.

We plan experiments that include a time dimension as well (i.e. timestamps of the texts), to see if the properties of the text on the Internet are already changing, due to the use of modern language tools.

Overall, we did not find conclusive proof of MT producing more uniformly surprising texts yet. Even if the future planned experiments end inconclusively, the objectives **O3** and **O4**, which encompass most of the future work, can stand on their own. The proposed novel decoding algorithms and training objectives could improve other aspects of language processing and generation as well.

5.2 **O3: Decoding algorithms**

So far, we have only considered problems brought by the beam search decoding and we have not explored sampling algorithms used in current LLMs.

Figuring out their influence on the properties of the text generated by LLMs will be the main part of the remaining work in this objective.

5.3 **O4: Training objectives**

The main focus of our future work will be on global training objectives. Such objectives could potentially improve the overall quality of natural language generation, aside from addressing the goals of our work. Thus, even if we fail to identify concrete issues in **O2**, developing these objectives can stand on its own. Global planning is linked to intrinsic uncertainty in MT, i.e. the notion that a single source sentence has multiple possible translations. This fact is not modeled in the current objective of the MT – each source example has only one target sentence probability distribution. This problem is more pronounced in globally operating training objectives since the uncertainty rises with increasing the distance from the prefix generated so far.

Our CPC objective addresses this by predicting future internal representation instead of a single token. However, even the representations used currently are only single-point estimates in the multi-dimensional space of the model. We will look into incorporating the uncertainty in these embeddings as well, similar to Kesiraju et al. (2020), which could help to address the inherent ambiguity in generating translations.

We will explore the use of segment-level training objectives and the effects of teacher forcing and its possible alternatives. We evaluate the ability of the alternative objectives to follow multiple possible paths in the translation using a multi-reference dataset (Bojar et al., 2013).

In the later stages, we plan to focus on LLMs with their specifics, like RLHF (which can also be considered a segment-level objective).

We might also move even further from traditional language modeling. An interesting approach to increase diversity and do away with the “single correct past” problem of auto-regressive, teacher-forced architectures is the use of diffusion models (Singh et al., 2023; Li et al., 2022; Lin et al., 2023).

5.4 **O5: Real world use-cases**

The more specific use cases (aside from saving the world from uniformity by fixing NLP) are closely related to the types of text we identify as problematic. So far, given both our intuition and experimental results in Section 4.2, literary translation seems like a promising testbed for our approach.

Limitations

This work has several limitations that should be taken into account when interpreting the preliminary results. First, many of the proposed diversity and surprisal-based measures are highly sensitive to methodological choices, including the language model used to estimate surprisal, tokenization schemes, normalization procedures, and dataset selection. As shown in our preliminary experiments, small changes in these factors can lead to qualitatively different conclusions, which limits the robustness and generalizability of individual findings. Second, our empirical evidence for excessive uniformity in NMT outputs remains mixed and is currently strongest only for highly unorthodox text types such as poetry, making it difficult to draw broad conclusions across domains and genres. Finally, while NMT serves as a useful testbed, the extent to which the observed phenomena and proposed solutions transfer to large-scale LLMs and real-world deployment scenarios remains an open question.

Acknowledgements

This work was partially supported by SVV project number 260 821, by Czech Ministry of Education, Youth and Sports (grant MŠMT OP JAK Mezisektorová spolupráce CZ.02.01.01/00/23_020/0008518) and by National Recovery Plan funded project MPO 60273/24/21300/21000 CEDMO 2.0 NPO.

It has been using language resources and tools developed and/or stored and/or distributed by the LINDAT/CLARIAH-CZ project of the Ministry of Education, Youth and Sports of the Czech Republic (project LM2023062).

References

Joshua Ainslie, Santiago Ontanon, Chris Alberti, Vaclav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. 2020. [ETC: Encoding long and structured inputs in transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 268–284, Online. Association for Computational Linguistics.

Chantal Amrhein and Rico Sennrich. 2022. [Identifying weaknesses in machine translation metrics through minimum Bayes risk decoding: A case study for COMET](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International*

Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 1125–1141, Online only. Association for Computational Linguistics.

- Matthew Aylett and Alice Turk. 2004. [The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech](#). *Language and Speech*, 47(1):31–56. PMID: 15298329.
- Alan Bell, Daniel Jurafsky, Eric Fosler-Lussier, Cynthia Girand, Michelle Gregory, and Daniel Gildea. 2003. [Effects of disfluencies, predictability, and utterance position on word form variation in English conversation](#). *The Journal of the Acoustical Society of America*, 113(2):1001–1024.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Ondřej Bojar, Matouš Macháček, Aleš Tamchyna, and Daniel Zeman. 2013. [Many czech references for 50 sentences selected from WMT11 data](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Mojca Brglez and Spela Vintar. 2022. [Lexical diversity in statistical and neural machine translation](#). *Inf.*, 13:93.
- Martin Briesch, Dominik Sobania, and Franz Rothlauf. 2023. [Large language models suffer from their own output: An analysis of the self-consuming training loop](#).
- Bryan Eikema and Wilker Aziz. 2020. [Is MAP decoding all you need? the inadequacy of the mode in neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4506–4520, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Bryan Eikema and Wilker Aziz. 2022. [Sampling-based approximations to minimum Bayes risk decoding for neural machine translation](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10978–10993, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- August Fenk and Gertraud Fenk-Oczlon. 1980. [Konstanz im kurzzeitgedächtnis - konstanz im sprachlichen informationsfluß? Zeitschrift für experimentelle und angewandte Psychologie](#), 27:400–414.

- Patrick Fernandes, António Farinhas, Ricardo Rei, José G. C. de Souza, Perez Ogayo, Graham Neubig, and Andre Martins. 2022. [Quality-aware decoding for neural machine translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1396–1412, Seattle, United States. Association for Computational Linguistics.
- Irene Fernandez Monsalve, Stefan L. Frank, and Gabriella Vigliocco. 2012. [Lexical surprisal as a general predictor of reading time](#). In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 398–408, Avignon, France. Association for Computational Linguistics.
- Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. 2022. [High quality rather than high model probability: Minimum Bayes risk decoding with neural metrics](#). *Transactions of the Association for Computational Linguistics*, 10:811–825.
- Richard Futrell, Edward Gibson, Harry J. Tily, Idan Blank, Anastasia Vishnevetsky, Steven Piantadosi, and Evelina Fedorenko. 2018. [The natural stories corpus](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Dmitriy Genzel and Eugene Charniak. 2002. [Entropy rate constancy in text](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 199–206, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Vaibhava Goel and William J. Byrne. 2000. [Minimum bayes-risk automatic speech recognition](#). *Comput. Speech Lang.*, 14:115–135.
- Adam Goodkind and Klinton Bicknell. 2018. [Predictive power of word surprisal for reading times is a linear function of language model quality](#). In *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)*, pages 10–18, Salt Lake City, Utah. Association for Computational Linguistics.
- Kartik Goyal, Chris Dyer, and Taylor Berg-Kirkpatrick. 2019. [An empirical investigation of global and local normalization for recurrent neural sequence models using a continuous relaxation to beam search](#). 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 1724–1733, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alex Graves. 2012. [Sequence transduction with recurrent neural networks](#).
- Yanzhu Guo, Guokan Shang, Michalis Vazirgiannis, and Chloé Clavel. 2023. [The curious decline of linguistic diversity: Training language models on synthetic text](#).
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Dávid Javorský, Ondřej Bojar, and François Yvon. 2023. [Assessing word importance using models trained for semantic tasks](#).
- Josef Jon and Ondřej Bojar. 2023. [Breeding machine translations: Evolutionary approach to survive and thrive in the world of automated evaluation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2191–2212, Toronto, Canada. Association for Computational Linguistics.
- Josef Jon and Ondřej Bojar. 2024. [Gaatme: A genetic algorithm for adversarial translation metrics evaluation](#). In *LREC-COLING 2024: The 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*. In print.
- Josef Jon, Martin Popel, and Ondřej Bojar. 2022. [CUNI-bergamot submission at WMT22 general translation task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 280–289, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Josef Jon, Martin Popel, and Ondřej Bojar. 2023. [CUNI at WMT23 general translation task: MT and a genetic algorithm](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 119–127, Singapore. Association for Computational Linguistics.
- Junczys-Dowmunt. 2020. [Marian :: Is MT really lexically less diverse than human translation? — marian-nmt.github.io](#). <https://marian-nmt.github.io/2020/01/22/lexical-diversity.html>. [Accessed 22-12-2025].
- Katharina Kann, Sascha Rothe, and Katja Filippova. 2018. [Sentence-level fluency evaluation: References help, but can be spared!](#) In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 313–323, Brussels, Belgium. Association for Computational Linguistics.
- Santosh Kesiraju, Oldrich Plchot, Lukas Burget, and Suryakanth V. Gangashetty. 2020. [Learning document embeddings along with their uncertainties](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2319–2332.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.

- Svenja Kranich. 2014. [Translations as a locus of language contact](#).
- Roger Levy and T. Florian Jaeger. 2006. Speakers optimize information density through syntactic reduction. volume 19, pages 849–856.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B. Hashimoto. 2022. [Diffusion-lm improves controllable text generation](#).
- Zhenghao Lin, Yeyun Gong, Yelong Shen, Tong Wu, Zhihao Fan, Chen Lin, Nan Duan, and Weizhu Chen. 2023. [Text generation with diffusion language models: A pre-training approach with continuous paragraph denoise](#).
- Wenjie Lu, Leiying Zhou, Gongshen Liu, and Qunhai Zhang. 2020. [A mixed learning objective for neural machine translation](#). In *Proceedings of the 19th Chinese National Conference on Computational Linguistics*, pages 974–983, Haikou, China. Chinese Information Processing Society of China.
- Philip M McCarthy. 2005. *An assessment of the range and usefulness of lexical diversity measures and the potential of the measure of textual, lexical diversity (MTLD)*. Ph.D. thesis, The University of Memphis.
- Clara Meister, Ryan Cotterell, and Tim Vieira. 2020. [If beam search is the answer, what was the question?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2173–2185, Online. Association for Computational Linguistics.
- Clara Meister, Tiago Pimentel, Patrick Haller, Lena Jäger, Ryan Cotterell, and Roger Levy. 2021. [Revisiting the Uniform Information Density hypothesis](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 963–980, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mathias Müller and Rico Sennrich. 2021. [Understanding the properties of minimum Bayes risk decoding in neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 259–272, Online. Association for Computational Linguistics.
- Khanh Nguyen and Hal Daumé III. 2019. [Global Voices: Crossing borders in automatic news summarization](#). In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*, pages 90–97, Hong Kong, China. Association for Computational Linguistics.
- Franz Josef Och and Hermann Ney. 2004. [The alignment template approach to statistical machine translation](#). *Computational Linguistics*, 30(4):417–449.
- Byung-Doh Oh and William Schuler. 2022. [Why does surprisal from larger transformer-based language models provide a poorer fit to human reading times?](#)
- 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.
- Myle Ott, Michael Auli, David Grangier, and Marc’ Aurelio Ranzato. 2018. [Analyzing uncertainty in neural machine translation](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#).
- Koyena Pal, Jiuding Sun, Andrew Yuan, Byron Wallace, and David Bau. 2023. [Future lens: Anticipating subsequent tokens from a single hidden state](#). In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 548–560, Singapore. Association for Computational Linguistics.
- Adam Pauls and Dan Klein. 2012. [Large-scale syntactic language modeling with treelets](#). In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 959–968, Jeju Island, Korea. Association for Computational Linguistics.
- Marc’ Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#).
- Ricardo Rei, Marcos Treviso, Nuno M. Guerreiro, Chrysoula Zerva, Ana C Farinha, Christine Maroti, José G. C. de Souza, Taisiya Glushkova, Duarte Alves, Luisa Coheur, Alon Lavie, and André F. T. Martins. 2022. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Cory Shain, Clara Meister, Tiago Pimentel, Ryan Cotterell, and Roger Levy. 2024. [Large-scale evidence for logarithmic effects of word predictability on reading time](#). *Proceedings of the National Academy of Sciences*, 121(10):e2307876121.
- Mukul Singh, José Cambronero, Sumit Gulwani, Vu Le, Carina Negreanu, and Gust Verbruggen. 2023. [Code-fusion: A pre-trained diffusion model for code generation](#).
- Nathaniel J Smith and Roger Levy. 2013. The effect of word predictability on reading time is logarithmic. *Cognition*, 128(3):302–319.

- Felix Stahlberg and Bill Byrne. 2019. [On NMT search errors and model errors: Cat got your tongue?](#) 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 3356–3362, Hong Kong, China. Association for Computational Linguistics.
- Felix Stahlberg, Iliia Kulikov, and Shankar Kumar. 2022. [Uncertainty determines the adequacy of the mode and the tractability of decoding in sequence-to-sequence models.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8634–8645, Dublin, Ireland. Association for Computational Linguistics.
- Felix Stahlberg and Shankar Kumar. 2022. [Jam or cream first? modeling ambiguity in neural machine translation with SCONES.](#) In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4950–4961, Seattle, United States. Association for Computational Linguistics.
- Mitchell Stern, Noam Shazeer, and Jakob Uszkoreit. 2018. [Blockwise parallel decoding for deep autoregressive models.](#)
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2022. [Learning to summarize from human feedback.](#)
- Antonio Toral. 2019. [Post-editeese: an exacerbated translationese.](#) In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 273–281, Dublin, Ireland. European Association for Machine Translation.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2019. [Representation learning with contrastive predictive coding.](#)
- Eva Vanmassenhove, Dimitar Shterionov, and Matthew Gwilliam. 2021. [Machine translationese: Effects of algorithmic bias on linguistic complexity in machine translation.](#) In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2203–2213, Online. Association for Computational Linguistics.
- Eva Vanmassenhove, Dimitar Shterionov, and Andy Way. 2019. [Lost in translation: Loss and decay of linguistic richness in machine translation.](#) In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 222–232, Dublin, Ireland. European Association for Machine Translation.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments.](#) *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei, Clara Meister, and Ryan Cotterell. 2021. [A cognitive regularizer for language modeling.](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5191–5202, Online. Association for Computational Linguistics.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books.](#) In *The IEEE International Conference on Computer Vision (ICCV)*.
- Vilém Zouhar and Ondřej Bojar. 2024. [Quality and quantity of machine translation references for automated metrics.](#)
- Vilém Zouhar, Věra Kloudová, Martin Popel, and Ondřej Bojar. 2023. [Evaluating optimal reference translations.](#)

A Progress in O1: Information distribution in language

Surprisal theory, introduced by [Hale \(2001\)](#), relates cognitive effort to the surprisal value of words, suggesting that the effort to comprehend a word increases with its unpredictability given the context. The surprisal of an element u_n in an utterance \mathbf{u} is defined as $s(u_n) = -\log p(u_n|\mathbf{u}_{<n})$. The theory ([Hale, 2001](#)) proposes that cognitive effort is proportional to surprisal:

$$\text{Effort}(u_n) \propto s(u_n)$$

Choosing reading time and linguistic acceptability ratings as proxies for the effort, we compared surprisal estimates computed by multiple language models across various datasets ([Smith and Levy, 2013](#); [Futrell et al., 2018](#); [Warstadt et al., 2019](#)) and we successfully confirmed this hypothesis on the word-level, as did many works before us ([Fernandez Monsalve et al., 2012](#); [Goodkind and Bicknell, 2018](#); [Oh and Schuler, 2022, 2023](#)).

Applying this concept to a longer sequence, like a sentence, leads to an unexpected conclusion: If we calculate the sentence’s surprisal as the sum of surprisals of its individual words, and if this aggregate surprisal predicts the effort needed to process the sentence, then any way of distributing the information across the utterance is the same in terms of the effort needed for comprehension.

The Uniform Information Density (UID) theory, proposed by [Levy and Jaeger \(2006\)](#), accounts for

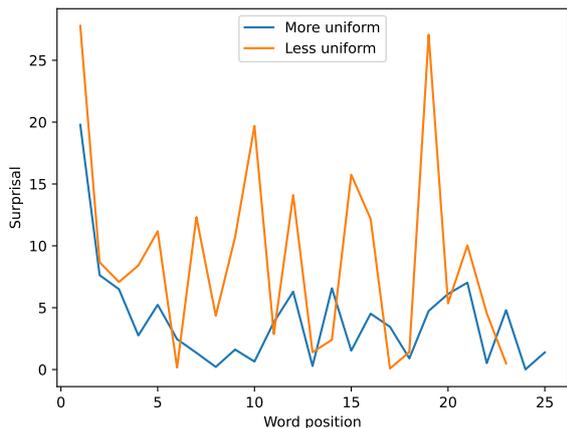


Figure 3: Surprisal behavior for the two examples sentences, measured by GPT-2 model.

this unintuitive conclusion by proposing a super-linear relationship between surprisal levels and cognitive effort, factoring in utterance length N , suggesting sentences with evenly distributed surprisal are easier to comprehend:

$$\text{Effort}(\mathbf{u}) \propto \sum_{n=1}^N s(u_n)^k + c \cdot N, k > 1$$

To illustrate the intuitive concept of surprisal uniformity, consider these two sentences:

- A) More uniform:** *"When she got home after a long day at work, she decided to relax by reading her favorite novel and having a cup of tea."*
- B) Less uniform:** *"London's annual festival was filled with activities, food stands, windsurfing, and drinks, but the sudden unveiling of a Yetti statue caught everyone's attention."*

Most people would consider the second sentence as more surprising, as some of the words feel unexpected. We show the surprisal profiles of both sentences in Figure 3. Indeed, we can see that the profile of the second sentence (orange) looks less uniform.

To operationalize our concept of uniformity and anchor it into the real world, we carried out assessments of various surprisal distribution uniformity measures for their correlation with human acceptability ratings and reading times, extending the work of Meister et al. (2021). We explore measures like Local Variance (LV), Coefficient of Variation (CV), Global Variance (GV), Gini coefficient, and Super-linear Relationship (SL), and super-linear syntactic log-odds ratio (SLOR, Kann et al., 2018; Pauls and Klein, 2012):

- $\text{LV}(\mathbf{u}) = \frac{1}{N-1} \sum_{n=2}^N (s(u_n) - s(u_{n-1}))^2$
- $\text{CV}(\mathbf{u}) = \frac{\sigma(\mathbf{u})}{\mu(\mathbf{u})}$
- $\text{GV}(\mathbf{u}) = \frac{1}{N} \sum_{n=1}^N (s(u_n) - \mu(\text{corpus}))^2$
- $\text{SL}(\mathbf{u}) = \frac{1}{N} \sum_{n=1}^N s(u_n)^k \quad (k > 1)$
- $\text{SLOR}(\mathbf{u}) = \frac{1}{N} \sum_{n=1}^N s(u_n)^k - s_u(u_n)^k \quad (k > 1)$

Function s denotes surprisal in of a word in context, s_n is a unigram, context-free surprisal. We aimed to predict a sentence's linguistic acceptability and reading times using these surprisal distribution uniformity measures, employing statistical and machine learning methods like Pearson's r, linear regression (LR), Support Vector Machines (SVM), Multi-layer perceptron (MLP), Generalized linear models (GLM) and Linear mixed-effect models (LME). We sought to mostly replicate the results of Meister et al. (2021) and to add novel evaluation methods. We succeeded in obtaining similar results under the exact same conditions, i.e. some evidence to support the existence of a super-linear (SL) relationship. However, we have also learned that the outcome is very sensitive to methodological choices, like the dataset, language model used to calculate the surprisal, preprocessing and filtering choices and evaluation methods. In fact, the SL relationship for sentence-level reading times was refuted by the most recent and most extensive study so far (Shain et al., 2024). This study presents strong evidence in favor of a simple linear relationship between word-level surprisals and effort (the team also included the authors of Meister et al. (2021)).

We present the results for linguistic acceptability for the SLOR measure in Figure 4. We see that the measure correlates better with acceptability ratings for $k > 1$, and it is slightly more predictive in LR and GLM models. According to SLOR, there might be some preference for sentences with more uniform surprisal distribution.

B Progress in O1 and O2: MT and uniformity of surprisal

We use the measures introduced in the previous section to confirm one of our initial hypotheses: *Is NMT producing translations that are more uniform in terms of surprisal distribution than a human?* This hypothesis posits that source sentences characterized by highly uneven surprisal distributions would exhibit more uniformity upon translation by

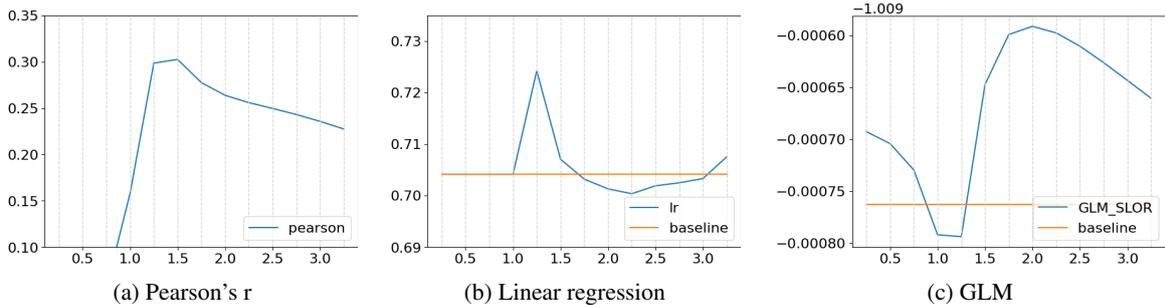


Figure 4: Behavior of SLOR measure, depending on power k used in the calculation.

MT systems, a phenomenon not expected to occur with human translations.

We measured this across multiple datasets: For English to French, we use the Books corpus (Zhu et al., 2015) (*books*), Global Voices (Nguyen and Daumé III, 2019) (*global*), Newstest2014 (Bojar et al., 2014) (*wmt*), and a poem by Oscar Wilde translated into French by Jean Guiloineau (*wilde*). For experiments in the English-Czech direction, we draw upon the dataset provided by Zouhar and Bojar (2024); Zouhar et al. (2023) (*ORT*).

dataset	measure	HT	MT1	MT2
books	LV^2	0.39	0.58	0.42
	CV	0.42	0.51	0.50
	GV^2	0.46	0.69	0.54
wmt	LV^2	0.43	0.54	0.58
	CV	0.49	0.55	0.57
	GV^2	0.46	0.57	0.64
global	LV^2	0.69	0.73	0.78
	CV	0.65	0.70	0.63
	GV^2	0.72	0.74	0.80
global_doc	LV^2	0.72	0.79	0.83
	CV	0.68	0.81	0.82
	GV^2	0.76	0.83	0.86
wilde	LV^2	0.16	0.40	0.53
	CV	0.07	0.39	0.54
	GV^2	0.16	0.40	0.53

Table 1: Pearson’s r for sentence-level surprisal uniformity of measurements between source and either HT, MT1 or MT2.

Table 1 reveals that machine translations (MT_x) exhibit a better correlation with the source text’s surprisal distribution than human translations (HT) for all measured indices: LV^2 (local variance squared), CV (coefficient of variation), and GV^2 (global variance squared). Considering that human translators might spread surprisal over larger text units, we extended our analysis to document level in the *global_doc* dataset, treating each document as a single sequence of tokens for the purposes of surprisal estimation. Yet, the results did not support

our hypothesis.

However, the absolute values of the uniformity measures also indicated that MT is generally (with some exceptions, depending on the measure and the dataset) as uniform or less uniform than HT. This contradicts our initial hypothesis that MT will be more uniform in surprisal. We hypothesized that this discrepancy might stem from errors in MT: if the MT system translates the input with some obvious mistakes, then these mistakes might be very surprising given the rest of the sentence. We used reference-free COMET (wmt22-cometkiwi-da Rei et al., 2022) scores to estimate the translation quality of the MT.

Figure 5 illustrates the LV^2 measure’s trends for instances where the machine translation (MT) COMET score surpasses a certain threshold (displayed on the x -axis). Uniformity between HT (green) and MT (two systems, blue and orange) remains steady across datasets, except for the *wilde* dataset, which exhibits greater HT unevenness in high-scoring translations. This observation could imply that MT achieves higher surprisal uniformity than HT in highly creative content like poetry, when inaccurately translated segments are excluded from the evaluation.

We have also experimented with another dataset, *ORT*, which contains an English source sentence and four high-quality Czech translations. We compared surprisal distribution uniformity between the hypotheses and two MT engines.

The results are presented in Figure 6. We see that MT1 usually scores as the most uniform, while MT2 is among the least uniform translations, showing large variance among different MT systems. The results, again, contradict our initial hypothesis, that MT is inherently more uniform than HT – it depends on both the human translator and MT system used. We again set a COMET threshold to filter out examples with low-quality MT. We

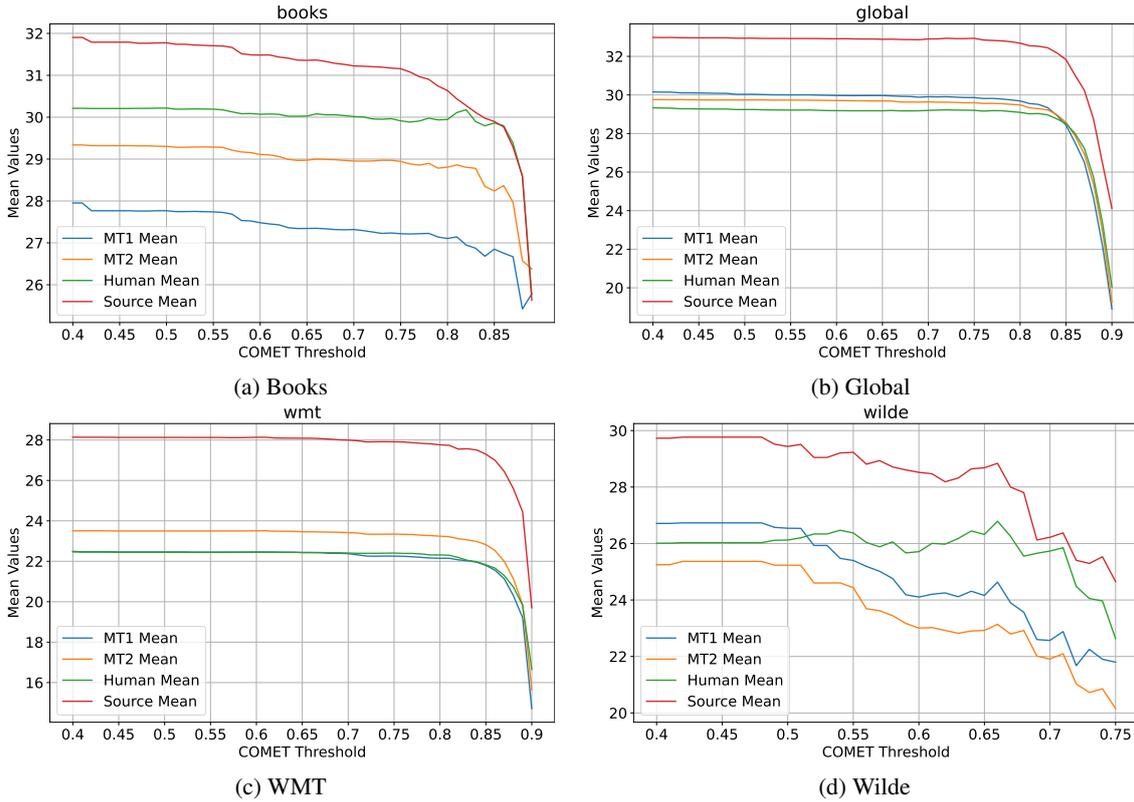


Figure 5: Relationship between COMET scores of the MT and the LV^2 measure. As a proxy of translation quality, we use COMET score threshold to filter out low-quality translations. Higher values of LV^2 mean more diverse surprisal distribution.

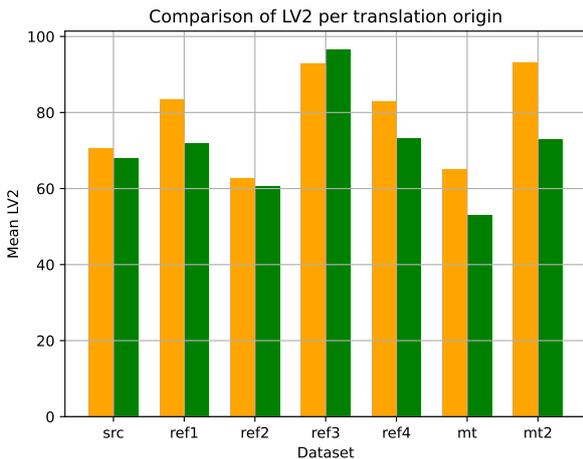


Figure 6: Difference of LV^2 scores between all (orange) and high-quality (green) MT translations. Higher scores indicate less uniform text.

see that surprisal diversity is lower in high-quality translations (green), especially for MT2, and the third human reference (ref3) becomes the most diverse. MT2 is as diverse as the human references, with ref3 being an exception. Overall, we do not have reliable proof that MT produces texts that are more uniform in surprisal distribution than humans yet. Either our hypothesis is false, or our measurement methodology is flawed. One possible reason could be that the LMs we used to estimate the surprisals are trained on human text, not on MT outputs so it overestimates surprisal of some phenomena in MT. We plan further experiments to improve our methodology and extend the analysis to more datasets.

C Progress in O3: Alternative decoding algorithms

We proposed a genetic algorithm to modify and rerank translation hypotheses. Illustration of the algorithm is presented in Figure 7. First, an initial set of candidate sentences is produced by an NMT model (e.g. via sampling or beam search). Candidates are then evaluated using a *fitness function*:

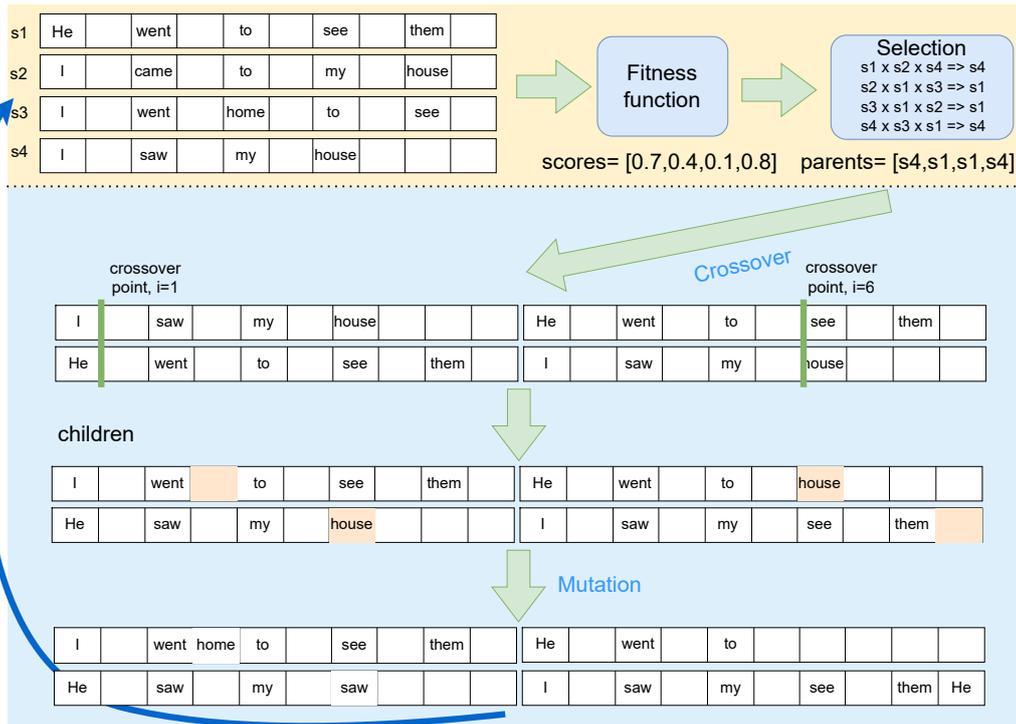


Figure 7: One iteration of the GA algorithm for a population of 4 individuals. The steps with the yellow background are equivalent to simple reranking, the steps with the blue background introduce the operations of the genetic algorithm.

a weighted sum of scores from various MT metrics, possibly using MBR decoding with the initial translation candidates as pseudo-references. The scores are used by the *selection method*, usually with some level of stochasticity (e.g. tournament selection), to choose the subset of translation candidates for as the parents for the next iteration of the genetic algorithm.

The selected parents undergo a *crossover* operation, where two translation hypotheses are split at a random word index and their segments swapped. Subsequently, the *mutation* operation randomly alters the candidates by removing, adding, or replacing tokens, drawing new token choices from the set of words in the initial candidates, dictionary, or the target language wordlist. These candidates are then scored again by the fitness function and the whole process is repeated. After a given number of iterations, the best-scoring candidate is selected as the final translation. The final translation always scores better or at least as well as the best translation from the initial candidates in the optimized metric.

Our findings indicate that employing a single MT metric in the fitness function tends to lead to overfitting, producing translations that achieve high

scores according to that metric but contain serious translation errors. To explore this behavior, we set aside a designated *held-out metric*, which is not used during the GA phase, to evaluate the translations generated by the algorithm. The results of this experiment can be interpreted as a rudimentary indicator of the robustness of the metric used in the fitness function: if the held-out metric scores are lower after GA than the initial hypotheses’ scores, i.e. if the GA process decreases the held-out scores, the metric is susceptible to adversarial examples and overfitting. In Table 2, we show for how many instances within our test set this decrease in held-out metric scores occurs. CMT20 and QE20 refer to *wmt20-comet-da* and *wmt20-comet-qe-da*. These results indicate the ease with which metrics can be misled to favor adversarially crafted translations that exploit their weaknesses, biases, and blind spots.

In Jon and Bojar (2024), we exploit this insight to create adversarial test sets for specific metrics, with a slight modification of the approach. Instead of using the held-out metric ex-post, to find the examples where the scores are worse, the held-out metric is directly incorporated into the fitness function, with a small negative weight. This change was

made to actively steer the process toward creating adversarial translations. We select random words for possible mutations from an English wordlist. See Table 4 in the appendix for examples of such adversarial translations for multiple metrics. Neural metrics show problems that are typical for them, like insensitivity to changes in named entities (probably due to embedding similarity of rare named entities) [Amrhein and Sennrich, 2022](#). BLEURT specifically seems to overwhelmingly prefer words that do not exist at all in the target language (they are part of the rare noise in the English wordlist).

O	$O_{init} + m_o < O_{ga}$	$\dots \wedge H_{init} > H_{ga} + m_h$
CMT20	128 (85%)	57 (38%)
QE20	148 (99%)	142 (95%)
BLEU	150 (100%)	113 (75%)

Table 2: Number of examples which improved in optimization metric after GA (2nd column) and at the same time deteriorated in held-out metric (3rd column). From [Jon and Bojar \(2023\)](#).

Fitness	+	-	=
BLEU	22%/1%	29%/7%	49%/92%
CHRf	13%/1%	69%/65%	18%/33%
CMT20	54%/23%	39%/32%	7%/45%
CMT20+QE20+BLEU	62%/43%	35%/35%	3%/23%

Table 3: Percentage of examples where the held-out score (UniTE) improves (+), degrades (-), or doesn’t change (=) for GA compared to best log-prob (before slash) or MBR reranking (after slash). Bold results are those where the held-out scores improve for more examples than deteriorate. From [Jon and Bojar \(2023\)](#).

On the other hand, we have found that the robustness can be improved by combining multiple metrics in the fitness function (we use a weighted sum of the values). In Table 3 we compare post-GA translations with the best (in terms of model log-prob) initial MT translation and the best translation after MBR reranking with the same metrics (i.e. after reranking with the fitness function used in the GA). We see that for the combination of CMT20, QE20, and BLEU beats or draws with the MBR reranking in most of the examples, while the standalone metrics (e.g. only CHRf) are performing worse, harming more examples than they improve. An illustration of behavior of fitness and held-out scores during GA in both the positive case (improved held-out score) and negative case (decreased held-out score) is shown in Figure 8 in the appendix.

These results were further extended and con-

firmed in our WMT23 submission ([Jon et al., 2023](#)), where we show that newer versions of the COMET metrics are more robust, but combining multiple metrics is still beneficial. We have used GA to modify the outputs of two other submissions in English to Czech and Czech to Ukrainian tracks. In both cases, the modified outputs were scored slightly better by the human evaluators, although the difference was not significant. We expect more improvement could be obtained by tuning the parameters (mutation or crossover rate, number of generations, ways to combine metrics in the fitness function, ...), which we did not do due to the computationally expensive nature of the process.

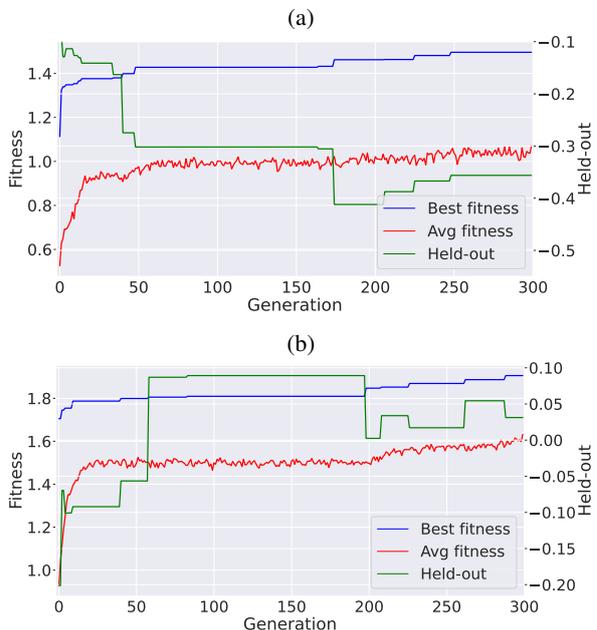


Figure 8: Behavior of best and population-average fitness, compared to held-out metric score H of the best solution during a run of GA for two selected examples. Held-out score H (UniTE) does not correlate well with the fitness metric (CMT20+QE+BLEU) and the GA is detrimental from the point of view of H in Example a). In Example b), H behaves similarly to the fitness function and the final held-out score is better than that of the best initial candidate. From [Jon and Bojar \(2023\)](#).

D Progress in O4: Alternative training objective

The concrete implementation is of our auxiliary CPC objective is through an InfoNCE loss, which maximizes the mutual information between an internal representation (in practice, we used embeddings in the last layer of the Transformer model) in the current step c_t and the future step c_{t+k} (positive

Metric	MT	post-GA	Ref	MT score	GA score
CMT22-QE	In the NHL, "France" caught 36 games, its save rate at 92.3%.	In yn , "Frederic" clocked up 36 ordain , with touchdown rate at 92.3	He has played 36 games in the NHL, where his save percentage is 92.3%.	0.6407	0.7679
	The 31-year-old full-back will be on the scoresheet and could soon be in goal.	The fullback will toilette on rotation and could get into goal soon fungo	The thirty-one-year-old Pilsen native will be on the bench and could soon be in goal.	0.6901	0.7242
CMT22	The highest ranked in the affair is Berbr , who no longer features in any of the football functions.	The highest profile in glave affair is Piute Denten who no longer longer figures stanno any football functions	The most senior figure in the affair is Berbr, who is no longer involved in any football function.	0.7947	0.8104
	Prince William, Duke of Cambridge , is wearing the same as Princes George and Louis shorts and a collared T-shirt.	Prince Pippo, Duke of Goldwyn , dressed the same as Princes Alexander and Louis in shorts and a T-shirt	Prince William, Duke of Cambridge, and Princes George and Louis are wearing shorts and a polo shirt.	0.7675	0.8402
CHRF	Interior has got respirators significantly cheaper than the Department of Health	Interior got respirators mushy Asch cheaper than the Ministry oie Ministry of Health natl . fur . LADT Goethe	The Ministry of the Interior got respirators much cheaper than the Ministry of Health	0.5342	0.8168
	PVO: medium-cold war, outdated; short range - good, modern, relatively good number.	unpaint : medium-cold war, obsolete; short range - good, modern, relatively Orth . enterable number. fugitively favorer POS SMDF R.A.A.F. pm . SM	SHORAD: medium - cold war, obsolete; short range - good, modern, relatively favorable number.	67.9	82.0
BLEU	The picture , which will serve as a Christmas card, was also posted by heir to the throne Prince Charles and wife Camilla.	knotty-leaved , which will be fat-shunning weeny-bopper Christmas card, was also posted by the heir to the throne, Prince Charles, dichlorodiphenyltrichloroethane duodenocholecystostomy cock-a-doodle-does his wife, Camilla.	The image, which will be used for the Christmas card, was also posted by the heir to the throne, Prince Charles, and his wife, Camilla.	34.3	68.6
BLEURT	By the time I got off my seat, it was gone.	Idun epicanthi got achenium tundun terebinthial off Ladakhi Morgenthaler gone.. ecliptically scholium mesonasal	By the time I got off the deer-stand, he was gone.	0.3965	0.8183
	His return to goal in the NHL eventually extended to more than two months.	succinimid Badajoz hootchie-kootchie cheongsam NHL taotai meromyarian Abyla Nadean vainer tenson months	In the end, his time away from the NHL was extended by more than two months.	0.5695	0.8510

Table 4: Examples from the adversarial test set. Superfluous words in the post-GA translation and words from before GA that are missing post-GA are in bold. From [Jon and Bojar \(2024\)](#).

example) and minimize it with respect to randomly selected representations from the same batch c_j (negative examples):

$$\mathcal{L}_N = -\mathbb{E}_C \left[\log \frac{f_k(c_{t+k}, c_t)}{\sum_{c_j \in C} f_k(c_j, c_t)} \right]$$

In practice, the function f_k can be implemented in many ways, we have used a log-bilinear model:

$$f_k(c_{t+k}, c_t) = \exp(c_{t+k}^T W_k c_t)$$

By minimizing this loss function, the mutual information between the representations is most preserved (derived in [van den Oord et al., 2019](#)). Note that for simplicity, we are using single token representation c_{t+k} as the target, but in practice, any embedding can be used, for example, a pooled, averaged embedding of multiple words, sentences, or paragraphs.