

Looking for a Needle in a Haystack: A Comprehensive Study of Hallucinations in Neural Machine Translation

Nuno M. Guerreiro^{1,2} Elena Voita⁴ André F. T. Martins^{1,2,3}

¹Instituto de Telecomunicações, Lisbon, Portugal

²Instituto Superior Técnico & LUMILIS (Lisbon ELLIS Unit), Lisbon, Portugal

³Unbabel, Lisbon, Portugal ⁴University of Edinburgh, Scotland

{nuno.s.guerreiro, andre.t.martins}@tecnico.ulisboa.pt lena-voita@hotmail.com

Abstract

Although the problem of hallucinations in neural machine translation (NMT) has received some attention, research on this highly pathological phenomenon lacks solid ground. Previous work has been limited in several ways: it often resorts to artificial settings where the problem is amplified, it disregards some (common) types of hallucinations, and it does not validate adequacy of detection heuristics. In this paper, we set foundations for the study of NMT hallucinations. First, we work in a *natural* setting, i.e., in-domain data without artificial noise neither in training nor in inference. Next, we annotate a dataset of over 3.4k sentences indicating different kinds of critical errors and hallucinations. Then, we turn to detection methods and both revisit methods used previously and propose using glass-box uncertainty-based detectors. Overall, we show that for preventive settings, (i) previously used methods are largely inadequate, (ii) sequence log-probability works best and performs on par with reference-based methods. Finally, we propose DEHALLUCINATOR, a simple method for alleviating hallucinations at test time which significantly reduces the hallucinatory rate.

1 Introduction

Neural machine translation (NMT) is becoming increasingly accurate (Vaswani et al., 2017; Akhbardeh et al., 2021), particularly in high resource language pairs where parallel data is abundant. However, even the best systems available today may generate *hallucinations*. These are extremely pathological translations that contain content that is unfaithful to the source sequence. Critically, a tiny fraction of these mistakes is all it takes to compromise user trust or safe deployment of NMT models in production.

Unfortunately, although the problem of hallucinations received some attention, research on this highly pathological phenomenon lacks solid

ground. First, previous work used multiple and often overlapping definitions and categories of hallucinations which makes it hard to draw connections between observations made in different works (Lee et al., 2018; Raunak et al., 2021; Zhou et al., 2021). Next, since hallucinations are extremely rare, previous work focused on settings in which the phenomenon is amplified, e.g. perturbing data either in training or at inference, or evaluating under domain shift (Lee et al., 2018; Raunak et al., 2021; Müller et al., 2020; Wang and Sennrich, 2020; Voita et al., 2021; Müller and Sennrich, 2021; Zhou et al., 2021). Critically, the analysis on these works mostly relied on the adequacy of the automatic hallucination detection methods they proposed. However, it is not immediate whether these methods translate well to unperturbed settings.

In this work, we set foundations for the study of NMT hallucinations. We take a step back from previous work and, instead of considering perturbed settings for which hallucinations are more frequent, we consider a *natural scenario* and face the actual problem of identifying a small fraction of hallucinations (a “needle”) in a large number of translated sentences (a “haystack”). Then, we provide a rigorous comparison among hallucination detection methods. Apart from analysing those proposed in previous work (e.g., heuristics based on anomalous encoder-decoder attention), we also propose to use simple model uncertainty measures as detectors. For each of these methods, we select examples marked as hallucinations, put them together, and gather human annotations. As a result, we introduce a corpus of 3415 structured annotations for different NMT pathologies and hallucinations. We use this corpus for analysis and show that, in preventive settings where high recall is desirable, previously proposed methods are mostly inadequate, and filtering according to standard sequence log-probability performs the best. In fact, it performs on par with the state-of-the-art COMET (Rei et al.,

Category	Source Sentence	Reference Translation	Hallucination
Oscillatory	Ist ein Kompromiss aufgrund des zugrundeliegenden Regelsystems unmöglich, so spricht man von Aporie.	The case where, based on the pertinent system of regulations a compromise is not possible, is referred to as Aporia.	Aporia is the name of aporia , which is the name of aporia.
Strongly Detached	Tickets für Busse und die U-Bahn ist zu teuer, vor allem in Stockholm.	Tickets for buses and the subway is too expensive, especially in Stockholm.	The hotel is located in the centre of Stockholm, close to the train station.
Fully Detached	Die Zimmer beziehen, die Fenster mit Aussicht öffnen, tief durchatmen, staunen.	Head up to the rooms, open up the windows and savour the view, breathe deeply, marvel.	The staff were very friendly and helpful.

Table 1: Examples of hallucination types. Hallucinated content is shown **shaded**.

2020a) which uses reference translation and thus cannot be used in most real-world on-the-fly applications. Surprisingly, its reference-free version COMET-QE (Rei et al., 2020b), which was shown to generally perform on par with COMET (Kocmi et al., 2021; Freitag et al., 2021), substantially fails to penalise the severity of hallucinations. Overall, methods targeting the phenomena are largely unfit, quality estimation systems fail, and sequence log-probability, i.e. a byproduct of generating a translation, turns out to be the best.

Apart from our analysis of detection methods, we propose DEHALLUCINATOR, a method for alleviating hallucinations at test time. At a high level, we first apply a lightweight hallucination detector and then, if a translation is flagged, we try to overwrite it with a better version. For this, we generate several MC-dropout hypotheses (Gal and Ghahramani, 2016), score them with some measure, and pick the highest-scoring translation as the final candidate. With this approach, the proportion of correct translations among the ones flagged by the detector increases from 33% to 85%, and the hallucinatory rate decreases threefold.

Overall, we show that (i) in preventive settings, previously proposed hallucination detectors are mostly inadequate; (ii) quality estimation techniques fail to distinguish hallucinations from less severe errors; (iii) sequence log-probability is the best hallucination detector and performs on par with reference-based COMET; and, (iv) our DEHALLUCINATOR significantly alleviates hallucinations at test time.

Additionally, we release our annotated dataset along with the model, training data, and code.¹

2 Taxonomy of Translation Pathologies

Choosing a good taxonomy is a compromise between simplicity (which minimizes annotation effort) and comprehensiveness. Thus, generic qual-

¹All these resources are available at <https://github.com/deep-spin/hallucinations-in-nmt>.

ity assessment taxonomies, such as MQM (Lommel et al., 2014), might be unfit or too complex when we focus only on critical errors and hallucinations. For hallucinations, in turn, previous work used multiple, often overlapping, definitions (Lee et al., 2018; Raunak et al., 2021; Zhou et al., 2021; Raunak et al., 2022). The taxonomy we build here is rather general: it covers categories considered previously (Lee et al., 2018; Raunak et al., 2021) and others not reported before. For a broader discussion on the taxonomy of hallucinations in NMT, and how it differs from other natural language generation tasks, refer to Ji et al. (2022).

2.1 Hallucinations

To distinguish hallucinations from other errors, we rely on the idea of detachment from the source sequence. From this perspective, other critical errors such as mistranslation of named entities are not considered as hallucinations. In Section 6 and Appendix D, we show that properties of hallucinations differ a lot from these other errors and thus our taxonomy is very reasonable.

Oscillatory hallucinations. These are inadequate translations that contain erroneous repetitions of words and phrases.

Largely fluent hallucinations. These are largely fluent translations that are unrelated to the content of the source sequence. Previous work assumed they always bear *no relation at all* to the source content (Lee et al., 2018; Raunak et al., 2021). However, we find that a large proportion of fluent hallucinations partially support the source. Therefore, we also consider severity of a hallucination and distinguish translations that are *fully* detached from those that are *strongly* (but not fully) detached.

Note that oscillatory hallucinations can also be either fully or only partially detached, but since these hallucinations are less frequent, in what follows we do not split them by severity. We show examples of these hallucination types in Table 1.

2.2 Translation errors

Undergeneration. These are incomplete translations that do not cover part of the source content. This problem is often studied in isolation (Koehn and Knowles, 2017; Stahlberg and Byrne, 2019; Kumar and Sarawagi, 2019). Undergenerations are sometimes considered as hallucinations (Lee et al., 2018) but we do not consider them so in our work.

Mistranslation of named entities. Appropriately translating named entities (e.g. names, dates, etc.) is also a known difficulty of NMT systems (Ugawa et al., 2018; Li et al., 2021; Hu et al., 2022). Note that for production systems, this error is rather critical. However, we do not consider it as an hallucination as it does not show detachment from the source but rather an incorrect attempt to translate part of its content.

Other errors. These are other incorrect translations that do not fit the categories above. They may include errors related to part of speech, word order, and others. For an extensive analysis on machine translation errors, refer to Vilar et al. (2006).

3 Hallucination Detection Methods

Approaches to hallucination detection generally aim to find low-quality translations that may also satisfy additional constraints. Previous work either relied only on quality filtering (Lee et al., 2018; Raunak et al., 2021; Müller and Sennrich, 2021), or only on heuristics (Berard et al., 2019), or a combination of the two (Raunak et al., 2021). We stick to this general form and consider different quality filters and heuristics.

3.1 Quality Filters

Quality filters come in two forms: reference-free and reference-based filters. The latter rely on reference translations, while the former do not.

Reference-free methods. We use the state-of-the-art COMET-QE (Rei et al., 2020b) for its superior performance compared to other metrics (Mathur et al., 2020; Freitag et al., 2021; Kocmi et al., 2021).

Reference-based methods. Previous work used adjusted BLEU or CHRF2 scores of less than 1% as a standalone criteria (Lee et al., 2018; Raunak et al., 2021; Müller and Sennrich, 2021; Yan et al., 2022). In this work, we analyse CHRF2 because it is more suitable for sentence-level evaluation. In

addition to this lexical metric, we also consider neural COMET (Rei et al., 2020a), a state-of-the-art reference-based metric (Kocmi et al., 2021).

Note that in real-world on-the-fly applications, detecting hallucinations is needed when references are not available. Thus, we use reference-based methods to estimate an upper bound for performance of the other methods.

3.2 Hallucination Detection Heuristics

3.2.1 Previously Used Heuristics

Binary-score Heuristics. These heuristics were used by Raunak et al. (2021) to detect oscillatory and fully detached hallucinations. Given a corpus of source-translation pairs, a translation is flagged as an hallucination if it is in the set of 1% lowest-quality translations, and if:

- **Top n -gram count (TNG).** The count of the top repeated n -gram in the translation is greater than the count of the top repeated source n -gram by at least t (in their work, $n = 4$ and $t = 2$);
- **Repeated targets (RT).** The translation is repeated for multiple unique source sentences.

Anomalous decoder-encoder attention. Attention patterns² in which most attention mass is concentrated on the source EOS token are often associated with a model ignoring the source and generating a hallucinatory translation (Lee et al., 2018; Berard et al., 2019; Raunak et al., 2021). We consider two different criteria targeted to find this pattern:

- **Attn-to-EOS:** the proportion of attention paid to the EOS source token;
- **Attn-ign-SRC:** the proportion of source words with a total incoming attention mass lower than 0.2. This was used as a data filtering criterion in Berard et al. (2019).

3.2.2 Uncertainty-Based Heuristics

Now we describe the uncertainty measures we propose to use as hallucination detectors. Previously, these were used to improve quality assessments (Fomicheva et al., 2020; Zerva et al., 2021).

²These patterns are respective to the average of the cross-attention heads of the decoder’s last layer.

Sequence log-probability (Seq-Logprob). For a trained model $P(y|x, \theta)$ and a generated translation y , Seq-Logprob (i.e., model confidence) is the *length-normalised* sequence log-probability:

$$\frac{1}{L} \sum_{k=1}^L \log P(y_k | y_{<k}, x, \theta). \quad (1)$$

We hypothesise that when hallucinating, a model is not confident.

Dissimilarity of MC hypotheses (MC-DSim). This method measures how the original hypothesis y disagrees with hypotheses $\{h_1, \dots, h_N\}$ generated in stochastic passes. For the same source sentence, we generate these new hypotheses using Monte Carlo (MC) Dropout (Gal and Ghahramani, 2016). Then we evaluate the average similarity:

$$\frac{1}{N} \sum_{i=1}^N \text{SIM}(h_i, y). \quad (2)$$

Different similarity measures can be used in place of SIM (e.g. METEOR (Banerjee and Lavie, 2005), BERTScore (Zhang et al., 2020), etc.). We follow previous work and use METEOR with $N = 10$ (Fomicheva et al., 2020; Zerva et al., 2021).

3.3 Trained Hallucination Detection Model

An exception from the general framework of hallucination detection is the work by Zhou et al. (2021) who *learn* to detect token-level hallucinations. Specifically, the authors create synthetic data where they randomly corrupt some tokens in a translation and reconstruct them with the BART model (Lewis et al., 2020). Then, the authors fine-tune a pretrained language model to identify the replaced tokens.

TokHal-Model. We evaluate the proportion of tokens that are predicted to be hallucinated and use this as a detection score.

3.4 Binary vs Continuous Scores

The methods above fall into two categories: *binary-score* and *continuous-score* heuristics. The former (only TNG and RT) output a value in $\{0, 1\}$, whereas the latter output a value in \mathbb{R} and the prediction is made depending on a chosen threshold.

4 Experimental Setting

Model. We use Transformer base (Vaswani et al., 2017) from fairseq (Ott et al., 2019).

Data. We use the WMT2018 DE-EN news translation data excluding Paracrawl (Bojar et al., 2018) – 5.8M sentence pairs. We randomly choose 2/3 of the dataset for training and use the remaining 1/3 as a held-out set for analysis. For validation, we use the *newstest2017* dataset.

Held-out Data Filtering. We are mainly interested in hallucinations produced for clean data. Since our held-out data comes from the WMT2018 training dataset and thus can be noisy, we filter it using Bicleaner, the filtering tool used in official releases of filtered ParaCrawl data (Sánchez-Cartagena et al., 2018; Ramírez-Sánchez et al., 2020; Kreutzer et al., 2022). Following previous work, we exclude examples with a score below 0.5 and end up with about 1.3M examples.

All details on preprocessing, hyperparameters and implementation can be found in Appendix A.

5 Hallucinations Dataset

To analyse the effectiveness of hallucination detection criteria, we pick a subset of examples that are likely to be hallucinations, and obtain fine-grained annotations from professional translators.

5.1 Data for Annotation

Our data selection is motivated by two goals: (i) find as many hallucinatory translations as possible – to analyse hallucinations, (ii) pick some translations from a long tail of hallucination detection predictions – to analyse the behaviour of these detection methods. Thus, we first pick 250 worst-scored samples for each heuristic and quality filter (including binary assignments obtained through TNG and RT). Next, we turn to a broader set of samples and consider translations whose scores fall below a chosen percentile for a given method, i.e. we consider long tails of the scores.³ For in-domain settings, previous work reported hallucinatory rates of 0.2 – 2%. However, these rates were either obtained on noisy and/or low-resource data or using weaker models. Therefore, in our cleaner in-domain setting with a stronger model, we expect the hallucination rate to lie in the lower end of the indicated range. Thus, we consider approximately 0.4% of the worst scores (which amounts to 5000 flagged translations for each criteria).⁴ From

³From now on, we refer to examples contained in a long tail of a method as “flagged” or “detected” by this method.

⁴The threshold for prediction is consistent with this percentile. For practical details, refer to Appendix G.

these, we sample 250 examples and add them to the dataset. In total, we end up with 3415 examples for annotation.⁵

5.2 Guidelines and Annotation

The annotation guidelines are developed according to the taxonomy defined in Section 2. All details on data collection can be found in Appendix B.

6 High-level Overview of the Dataset

6.1 General Statistics

Figure 1 gives a structured overview of dataset statistics. First, we see that while we picked translations that are likely to be pathological, 60% of the dataset consists of correct translations. This highlights that with the existing methods, finding poor translations reliably is still challenging. Next, note that most of the incorrect translations have translation errors that are not severe enough to be deemed hallucinations. This agrees with the view that hallucinations lie at the extreme end of the spectrum of MT pathologies (Raunak et al., 2021). Finally, the results of the annotation confirm that our data selection is very reasonable. Indeed, while previous work has reported hallucinatory rates of 0.2 – 2% in in-domain settings, we see that, for a reasonably numbered collection of examples, our hallucination rate is substantially higher (9%) – 294 hallucinations among the 3415 translations. In Figure 1, we also show the method-specific statistics of human annotation results for each heuristic and quality filter. Unsurprisingly, the long tails of each method display different characteristics. For example, almost all translations flagged by Attn-to-EOS are correct, whereas the proportion of good translations flagged by COMET-QE or Seq-Logprob is rather small. We will analyse this further in Section 7.

6.2 MT Errors vs Hallucinations

Now let us look separately at the sets of examples with hallucinations and other translation errors. Figure 2 shows the structures of these sets with respect to interactions of the different criteria.

We see that it is reasonable to consider hallucinations separately from other errors: patterns in Figures 2a and 2b are substantially different. For example, for translation errors, COMET-QE performs well, being on par with reference-aware

⁵Note that a sample originally obtained from the worst scores or sampled from the long tail of a given method may belong to the long tail of another method.

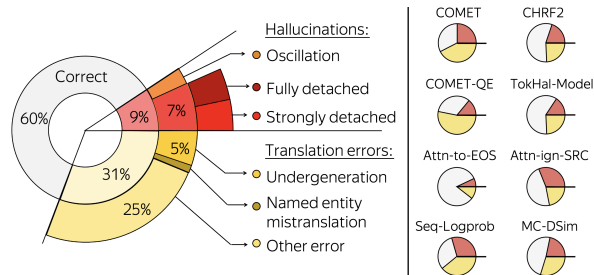


Figure 1: Overall (left) and method-specific (right) statistics of human annotation results. Method-specific statistics show the percentages of correct translations (grey), translation errors (yellow) and hallucinations (red) among the examples flagged by each method.

COMET (Figure 2b). However, for hallucinations, it identifies less than half the amount of examples identified not only by COMET, but even by simple uncertainty-based heuristics (Figure 2a). Furthermore, Figure 2 reveals a significant difference between interactions of different criteria for hallucinations and for other translations: hallucinations in our data are most often flagged by multiple criteria simultaneously (e.g. Seq-Logprob flags the majority of hallucinations detected with Attn-ign-SRC and MC-DSim), whereas most MT errors and correct translations are flagged by a single method. This difference in patterns supports our choice of taxonomy: properties of hallucinations are very different from those of other less severe errors.

7 Analysing Detection Criteria

In this section, we provide a comprehensive analysis of the performance of the heuristics and quality filters introduced in Section 3.

7.1 Quality Filters

Here we start by analysing reference-based methods, namely COMET and CHRF2, and then turn to the reference-free COMET-QE. Overall, our results show that reference information is helpful, while COMET-QE fails to penalise hallucinations.

Reference information helps detection. Figure 2a shows that leveraging reference information helps detecting hallucinations: COMET detects more hallucinations than any of the other methods. As expected, lexical-based CHRF2 is significantly worse than neural-based COMET. In fact, it lags behind several heuristics (e.g., Seq-Logprob).

We also explore whether previously proposed methods for detecting hallucinations under domain

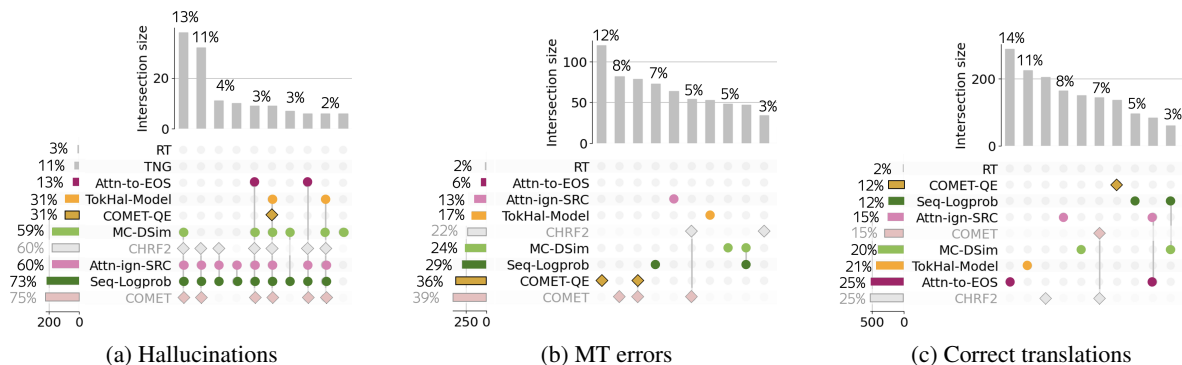


Figure 2: Structure of the sets of translations flagged by the considered methods. Horizontal bars show the proportion of examples flagged by each method among all translations of the considered category. Each vertical bar shows the size for the set of translations that are (i) flagged by all the methods marked in the corresponding column and (ii) not flagged by any of the rest; only the top intersections are shown. Quality filters are shown with diamond marks, and detection heuristics – with circles. Methods requiring reference translations are shaded.

Heuristic	Correct	MT Errors	Hallucinations		
			OSC	SD	FD
TNG	0	0	32	0	0
RT	18	19	2	1	7
All dataset	2048	1073	86	90	118

Table 2: Translations flagged by TNG and RT.

shift are appropriate in our cleaner in-domain setting. Specifically, we follow Müller and Senrich (2021) and consider translations with CHRF2 score lower than 1%. Strikingly, for our clean setting, this approach is inadequate: in the entire held-out set of 1.3M examples, it flagged only 2 translations. This suggests that methods suitable for noisy settings (e.g., domain shift) might not be applicable in settings where models are less likely to hallucinate.

COMET-QE fails to penalise hallucinations.

From Figure 1 (right) we see that, as expected, most of the translations flagged by COMET-QE are incorrect. However, the vast majority of them are not hallucinations. Indeed, Figure 2a shows that COMET-QE is one of the worst hallucination detectors, meaning that it fails to rank by the severity of a translation pathology. This supports the hypothesis made in previous work: since quality estimation models are mostly trained on data that lacks negative examples, they may be inadequate for evaluating poor translations (Takahashi et al., 2021; Sudoh et al., 2021).

Overall, among the considered quality filters, only COMET may be used as a hallucination detector. However, it is important to keep in mind that, in on-the-fly applications, detecting hallucina-

tions is needed when references are not available, rendering reference-based methods not applicable.

7.2 Detection Heuristics

The results in Section 7.1 leave a relevant gap: where do we turn to when references are not available? Is there any information, besides quality, that may help detecting hallucinations? Our evidence suggests that in preventive settings, previously proposed heuristics are mostly inadequate, and uncertainty may be the answer for the questions we pose.

Binary-score heuristics perform the worst. Table 2 shows the number of detected translations for Top n -gram count (TNG) and Repeated targets (RT). These heuristics are targeted to identify oscillatory and fully detached hallucinations, respectively. We see that while TNG obtains perfect precision, it fails to identify more than half of the oscillatory translations. RT, in turn, performs poorly across the board: only a few hallucinations are detected, and a significant proportion of flagged translations turn out to be correct. Moreover, Figure 2 shows that even if we join sets of translations detected by TNG and RT (as done in Raunak et al. (2021)), altogether we get *fewer hallucinations than almost any other considered method*. Thus, in preventive settings, these methods are highly inadequate.

Anomalous attention is mostly *not* hallucination.

Figures 1 and 2 show that behaviors of Attn-to-EOS and Attn-ign-SRC are significantly different. First, Attn-to-EOS is *not indicative of hallucinations*. Indeed, attention patterns in which most attention mass is concentrated on the EOS token

largely correspond to correct translations. On the other hand, Attn-ign-SRC performs well and is second only to uncertainty-based Seq-Logprob. Such a difference in performance is surprising: both methods are motivated by a common belief that if almost all the attention mass is concentrated on the source EOS token, a translation is likely to be a hallucination (Berard et al., 2019; Raunak et al., 2021). In fact, both methods were designed to identify this specific pattern. However, patterns identified with Attn-ign-SRC span from attention mass coming to various uninformative tokens (e.g., punctuation) to examples where attention is mostly diagonal (typically, these correspond to undergenerations). We show examples of such attention maps in Appendix C. Overall, the results highlight a disparity between what is *believed to indicate* hallucinations and what *actually indicates* them.

Note that while Attn-ign-SRC performs relatively well, it should be used with caution: attention-based heuristics rely on the assumption that attention patterns reflect model reasoning. This assumption is not reliable: although there is evidence that attention can play recognizable roles (Voita et al., 2018, 2019), a lot of work questions attention explainability (Wiegrefe and Pinter, 2019; Jain and Wallace, 2019; Serrano and Smith, 2019; Bastings and Filippova, 2020; Pruthi et al., 2020). Since hallucinations identified by Attn-ign-SRC are overwhelmingly contained among the ones identified by Seq-Logprob (Figure 2a), we recommend using the latter instead.

TokHal-Model is unfit for natural hallucinations.

Let us recall that the model used for the TokHal-Model scores was trained to identify replaced tokens in corrupted translations that are fluent and do not differ much from the original ones (Zhou et al., 2021). This means that during training, the model was unlikely to observe highly pathological translations that reflect the types of hallucinations produced by actual NMT systems. This raises several concerns when using this model in our setting. For example, it might incorrectly flag adequate tokens such as synonyms or paraphrases. What is more, since severely flawed examples are mostly out of distribution for the model, labels predicted for such translations may be unreasonable.

The results confirm our concerns: Figures 1 and 2a show that (i) the vast majority of translations flagged by TokHal-Model are correct and (ii) it is one of the worst hallucination detectors.

Model confidence may be all you need. Figure 2 shows that Seq-Logprob is *the best heuristic* and performs on par with reference-based COMET. This means that the less confident the model is, the more likely it is to generate an inadequate translation. This agrees with some observations made in previous work on quality estimation (Fomicheva et al., 2020). Interestingly, such performance of the method contrasts with its simplicity: Seq-Logprob scores are easily obtained as a by-product of generating a translation. This distinguishes the method from all the rest that require additional computation (e.g., corpus-level search for RT or generating multiple hypotheses for MC-DSim).

On another note, NMT models have been found to be miscalibrated (Kumar and Sarawagi, 2019). Thus, investigating the impact of recalibration methods – aimed at enhancing prediction reliability – on the detection quality of Seq-Logprob constitutes an interesting avenue for future research.

MC-DSim: hallucinations are mostly unstable.

The intuition behind this heuristic is simple: when faced with a source sentence for which a good translation is not immediate, the set of hypotheses the model “keeps at hand” may be very diverse. Indeed, MC-DSim performs relatively well and identifies a good proportion of hallucinations (Figures 1, 2). Thus, most hallucinations are *unstable*: dissimilarity of MC hypotheses helps to identify them.

Note that while MC-DSim is not the best choice for detection, later we will show that the intuition behind this method is very helpful for alleviating hallucinations at test time (Section 9).

7.3 Combining Heuristics and Quality Filters

Intersecting a set of translations obtained via a heuristic with the bottom scored translations according to a quality filter was introduced in Raunak et al. (2021). The motivation for this was to avoid incorrectly flagging good translations as hallucinations (e.g., without this intersection, RT flags good-quality paraphrases). Intuitively, this idea is very reasonable: hallucinations are indeed incorrect translations. However, implementing this in practice requires a good quality estimation model, specifically for ranking poor translations. Unfortunately, we showed in Section 7.1 that for this purpose, even the state-of-the-art COMET-QE is largely inadequate. This means that using such quality estimates may lead to filtering out a lot of hallucinations, which is not desirable in preven-

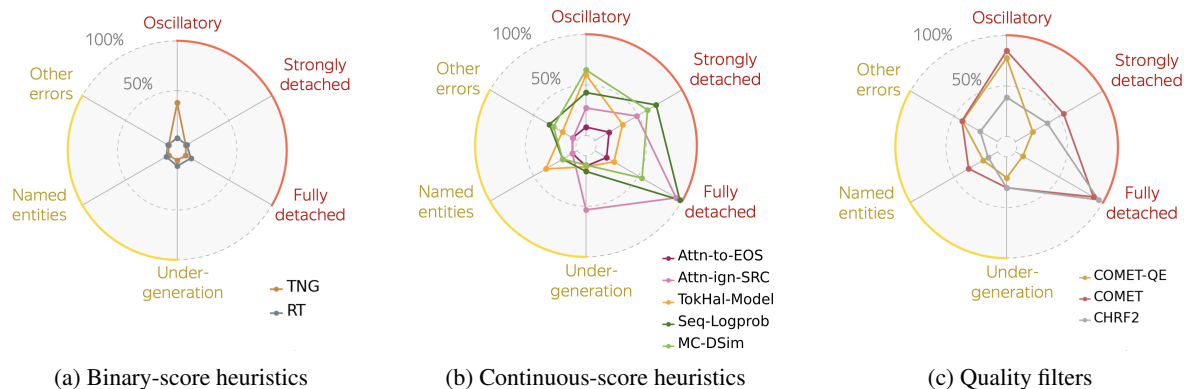


Figure 3: Distribution of the translations flagged by each method conditioned on each of the considered pathologies.

tive settings. Our results show that this is exactly the case: e.g., filtering with COMET-QE leads to losing nearly 80% of hallucinations detected by Seq-Logprob. All in all, such an intersection generally does more harm than good.

8 Analysing Hallucination Pathologies

In this section, we look at hallucination pathologies in isolation and show that the behavior of detection methods varies depending on the type of pathology. For example, for a given pathology, some methods may be specialised, whereas others may fail. For a similar analysis on other less severe translation errors, refer to Appendix D.

Fully detached hallucinations. Figure 3 shows that these hallucinations are easily detected by several methods, e.g. Seq-Logprob, Attn-ign-SRC, COMET, CHRF2. This is not surprising: intuitively, the most severe pathology should be the easiest to detect. However, COMET-QE *fails to identify almost all these hallucinations*. While COMET-based metrics are known to not penalise enough certain types of errors (e.g., discrepancies in numbers and named entities; see Amrhein and Sennrich (2022); Raunak et al. (2022)), such poor performance for completely inadequate translations is highly unexpected. This calls for further research on the behavior of quality estimation models.

On a more general note, previous work suggested that fully detached hallucinations emerge as exact copies of references from training data (Raunak et al., 2021). We validate this hypothesis and find the contrary: out of the 44 unique translations marked as fully detached from the source, only 4 are exact copies of references in the training data. Nevertheless, when looking at these sentences more closely, we see that they do contain large

substrings that are seen frequently during training. Therefore, fully detached hallucinations are indeed likely to be traced back to the training data, but they *emerge in non trivial ways and not necessarily as exact copies*. This can be seen as one more evidence that, when dealing with memorisation in language models, it is necessary to consider not just full copies in the training data but also near-duplicates (Lee et al., 2022).

Strongly detached hallucinations. As expected, this pathology is harder to detect than fully detached hallucinations (Figure 3). However, the trends are largely similar: for example, COMET-QE fails again, and Seq-Logprob performs best and outperforms even reference-based COMET.

Oscillatory hallucinations. The method specifically developed to detect this hallucination type (Top n -gram count, TNG) performs worse than most of the other methods. Among the rest, COMET performs best. Interestingly, in contrast to previous observations, COMET-QE performs well, being on par with COMET.

9 DEHALLUCINATOR: Overwriting Hallucinations at Test Time

In previous sections, we saw that hallucinations are more unstable than other translations: for them, generated MC-dropout hypotheses tend to vary greatly. This motivates us to look more closely into these hypotheses: are any of these translations *not* hallucinations? Answering this question not only gives insight into the inner workings of hallucinating NMT models, but also leads to an interesting practical application – overwriting hallucinations at inference time. This is of utmost importance for production systems where hallucinations have a deeply compromising effect on user trust.

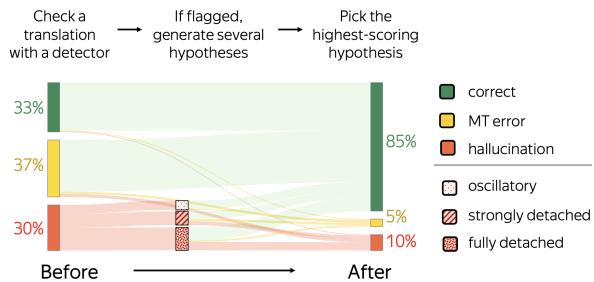


Figure 4: Our pipeline scheme along with results.

Whenever flagged, overwrite with better. Intuitively, our idea is similar to hybrid pipelines when a machine-generated translation is first passed to a quality estimation system and then, if needed, is corrected by human translators. In our case, we first apply a hallucination detector and then, if a translation is flagged, we try to overwrite it with a better translation (Figure 4). For this, we generate several MC-dropout hypotheses, score them with some measure, and pick the highest-scoring translation as a final candidate (in the spirit of reranking approaches (Shen et al., 2004; Lee et al., 2021; Fernandes et al., 2022; Freitag et al., 2022)).

The general pipeline above relies on the choice of a hallucination detector and a scoring measure. For the detector, we use the best of the analysed detectors, i.e. Seq-Logprob.⁶ For the scoring measure, a natural choice would be a quality estimation system: by construction, these systems are designed to score translations according to quality. However, as we saw earlier, even the state-of-the-art COMET-QE may fail (Section 7.1). Therefore, we compare two measures: COMET-QE and Seq-Logprob.

In this experiment, we randomly choose 200 translations from our dataset flagged by Seq-Logprob. For each, we generate 10 hypotheses with MC-dropout. Then, for the overwritten translations we gather annotations according to our guidelines (Section 5). The results are summarized in Figure 4. Although we were concerned about COMET-QE because of its low performance when ranking poor translations, we find that for choosing the best hypothesis, it is indeed appropriate and performs better than Seq-Logprob. We thus show results with COMET-QE scores in the main text and with Seq-Logprob in Appendix E.

Most hallucinations and errors become correct. Figure 4 shows that most hallucinations are over-

⁶In this experiment, we take the translations from our dataset and consider the percentiles defined in Section 5.

written with correct translations. This is surprising: in most cases, the model is not stuck in a hallucinatory mode and can generate good translations in a small vicinity of model parameters. In this sense, most hallucinations result from “bad luck” during generation and not profound model defect. Note that fully detached hallucinations are the hardest to improve. This makes sense as these are likely to be traced back to (near-)duplicates in the training data and, therefore, they do highlight model anomalies.

Note that in this pipeline, we overwrite not only hallucinations but also other errors and correct translations that were flagged by the detector. This means that our method needs to appropriately handle such translations. From Figure 4 we see a pleasant side-effect: our approach overwrites most of the errors with correct translations. Just as importantly, almost all originally correct translations remain correct. Overall, the proportion of correct translations among the ones flagged by the detector increases from 33% to 85%, and the hallucinatory rate decreases threefold. In Appendix F, we show several examples of overwritten hallucinations.

10 Conclusions

Dealing with hallucinations is difficult. First, we had to take a step back from previous work and refuse procedures that amplify the problem, as these hinder the behavior of models in their standard settings. After that, we notice that work on detection often relies on assumptions that remained unquestioned (e.g., generic quality measures, targeted heuristics, anomalous attention being suitable detectors). Through extensive experiments, we establish order in detection methods. Surprisingly, despite introduction of several methods specifically targeted for hallucinations, what works best has always been at our disposal: standard sequence log-probability. This suggests that characteristics innate to a model can have a lot of value. In fact, such characteristics are the backbone of our DEHALLUCINATOR, a lightweight approach that significantly alleviates hallucinations at test time. This leaves space for future research on model uncertainty, hallucination prevention, understanding where hallucinations come from, among others. For this, we release our corpus with structured annotations along with the model and its training data. Altogether, this allows us to say that we provide solid ground for future study of hallucinations in NMT.

Limitations

We highlight three main limitations in our work. First, although the foundation of our proposed taxonomy for hallucinations rests on the idea of detachment from the source content, we do not evaluate it quantitatively. Indeed, we cannot point whether the model that generated the hallucinations was indeed detached from the source sentence when generating them. Nevertheless, we can guarantee that the hallucinations in our dataset are translations that are detached from the source content according to professional translators. We consider that the quantitative analysis of the detachment to be out of the scope of this paper. Still, it constitutes an interesting line for future research on understanding hallucinations in NMT that may be facilitated with the release of our code, model and annotated dataset.

Second, while this paper comprehensively studies the phenomena of hallucinations in NMT for a high-resource language pair, experiments in more language pairs (including low-resource languages) are necessary to assess the broad validity of our claims. To keep our setup familiar to researchers and practitioners, we opted for a familiar language pair for which data is widely available. Moreover our choice also facilitated the data collection process as there is a large supply of professional translators for this language pair.

Lastly, instead of focusing on more recent NMT models that use large pretrained language models as their backbone, we focused on a Transformer base model. The reason for this choice is that we wanted to keep the setup simple, familiar, easy to reproduce, and computationally economical. Moreover, it was important for our work to have full control on the training and held-out data. Nevertheless, research on hallucinations on more recent and powerful NMT models is an exciting line of future work and we hope our work spurs that research.

Acknowledgments

This work is partially supported by the European Research Council (ERC StG DeepSPIN 758969, by the FCT through contract UIDB/50008/2020, by EU’s Horizon Europe (UTTER, HORIZON-CL4-2021-HUMAN-01-13, contract 101070631), and by the P2020 programs MAIA and Unbabel4EU (LISBOA-01-0247-FEDER-045909 and LISBOA01-0247-FEDER-042671). Lena is supported by the Facebook PhD Fellowship.

References

- Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydrin, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(WMT21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88, 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](#).
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Jasmijn Bastings and Katja Filippova. 2020. [The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?](#) In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, Online. Association for Computational Linguistics.
- Alexandre Berard, Ioan Calapodescu, and Claude Roux. 2019. [Naver labs Europe’s systems for the WMT19 machine translation robustness task](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 526–532, Florence, Italy. Association for Computational Linguistics.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Philipp Koehn, and Christof Monz. 2018. [Findings of the 2018 conference on machine translation \(WMT18\)](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels. Association for Computational Linguistics.
- Jacob Cohen. 1960. [A coefficient of agreement for nominal scales](#). *Educational and Psychological Measurement*, 20(1):37–46.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco

- Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Unsupervised cross-lingual representation learning at scale](#).
- Patrick Fernandes, António Farinhas, Ricardo Rei, José 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.
- Marina Fomicheva, Shuo Sun, Lisa Yankovskaya, Frédéric Blain, Francisco Guzmán, Mark Fishel, Nikolaos Aletras, Vishrav Chaudhary, and Lucia Specia. 2020. [Unsupervised quality estimation for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:539–555.
- 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.
- Markus Freitag, Ricardo Rei, Nitika Mathur, Chi-kiu Lo, Craig Stewart, George Foster, Alon Lavie, and Ondřej Bojar. 2021. [Results of the WMT21 metrics shared task: Evaluating metrics with expert-based human evaluations on TED and news domain](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 733–774, Online. Association for Computational Linguistics.
- Yarin Gal and Zoubin Ghahramani. 2016. [Dropout as a bayesian approximation: Representing model uncertainty in deep learning](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1050–1059, New York, New York, USA. PMLR.
- Junjie Hu, Hiroaki Hayashi, Kyunghyun Cho, and Graham Neubig. 2022. [DEEP: DENOISING ENTITY PRE-TRAINING FOR NEURAL MACHINE TRANSLATION](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1753–1766, Dublin, Ireland. Association for Computational Linguistics.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). 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 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. 2022. [Survey of hallucination in natural language generation](#).
- Tom Kocmi, Christian Federmann, Roman Grundkiewicz, Marcin Junczys-Dowmunt, Hitokazu Matsushita, and Arul Menezes. 2021. [To ship or not to ship: An extensive evaluation of automatic metrics for machine translation](#).
- 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.
- Julia Kreutzer, Isaac Caswell, Lisa Wang, Ahsan Wahab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmungkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iroko Orife, Kelechi Ogueji, Andre Niyongabo Rubungo, Toan Q. Nguyen, Mathias Müller, André Müller, Shamsuddeen Hassan Muhammad, Nanda Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, Mathias Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine Çabuk Ballı, Stella Biderman, Alessia Battisti, Ahmed Baruwa, Ankur Bapna, Pallavi Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. 2022. [Quality at a Glance: An Audit of Web-Crawled Multilingual Datasets](#). *Transactions of the Association for Computational Linguistics*, 10:50–72.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Aviral Kumar and Sunita Sarawagi. 2019. [Calibration of encoder decoder models for neural machine translation](#). *CoRR*, abs/1903.00802.
- Ann Lee, Michael Auli, and Marc’Aurelio Ranzato. 2021. [Discriminative reranking for 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 7250–7264, Online. Association for Computational Linguistics.
- Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fanfjang, and David Sussillo. 2018. [Hallucinations in neural machine translation](#).
- Katherine Lee, Daphne Ippolito, Andrew Nystrom, Chiyuan Zhang, Douglas Eck, Chris Callison-Burch, and Nicholas Carlini. 2022. [Deduplicating training data makes language models better](#). In *Proceedings of the 60th Annual Meeting of the Association for*

- Computational Linguistics (Volume 1: Long Papers)*, pages 8424–8445, Dublin, Ireland. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Panpan Li, Mengxiang Wang, and Jian Wang. 2021. [Named entity translation method based on machine translation lexicon](#). *Neural Comput. Appl.*, 33(9):3977–3985.
- Arle Lommel, Aljoscha Burchardt, and Hans Uszkor-eit. 2014. [Multidimensional quality metrics \(mqm\): A framework for declaring and describing translation quality metrics](#). *Tradumàtica: tecnologies de la traducció*, 0:455–463.
- Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondřej Bojar. 2020. [Results of the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 688–725, Online. Association for Computational Linguistics.
- Mathias Müller, Annette Rios, and Rico Sennrich. 2020. [Domain robustness in neural machine translation](#). In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 151–164, Virtual. Association for Machine Translation in the Americas.
- 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.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Danish Pruthi, Mansi Gupta, Bhuwan Dhingra, Graham Neubig, and Zachary C. Lipton. 2020. [Learning to deceive with attention-based explanations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4782–4793, Online. Association for Computational Linguistics.
- Gema Ramírez-Sánchez, Jaume Zaragoza-Bernabeu, Marta Bañón, and Sergio Ortiz-Rojas. 2020. [Bifixer and bicleaner: two open-source tools to clean your parallel data](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 291–298, Lisboa, Portugal. European Association for Machine Translation.
- Vikas Raunak, Arul Menezes, and Marcin Junczys-Dowmunt. 2021. [The curious case of hallucinations in neural machine translation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1172–1183, Online. Association for Computational Linguistics.
- Vikas Raunak, Matt Post, and Arul Menezes. 2022. [Salted: A framework for salient long-tail translation error detection](#).
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020a. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020b. [Unbabel’s participation in the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 911–920, Online. Association for Computational Linguistics.
- Víctor M. Sánchez-Cartagena, Marta Bañón, Sergio Ortiz-Rojas, and Gema Ramírez-Sánchez. 2018. [Prompsit’s submission to wmt 2018 parallel corpus filtering shared task](#). In *Proceedings of the Third Conference on Machine Translation, Volume 2: Shared Task Papers*, Brussels, Belgium. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Libin Shen, Anoop Sarkar, and Franz Josef Och. 2004. [Discriminative reranking for machine translation](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter*

- of the Association for Computational Linguistics: HLT-NAACL 2004, pages 177–184, Boston, Massachusetts, USA. Association for Computational Linguistics.
- 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.
- Katsuhito Sudoh, Kosuke Takahashi, and Satoshi Nakamura. 2021. [Is this translation error critical?: Classification-based human and automatic machine translation evaluation focusing on critical errors.](#) In *Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 46–55, Online. Association for Computational Linguistics.
- Kosuke Takahashi, Yoichi Ishibashi, Katsuhito Sudoh, and Satoshi Nakamura. 2021. [Multilingual machine translation evaluation metrics fine-tuned on pseudo-negative examples for wmt 2021 metrics task.](#) In *Proceedings of the Sixth Conference on Machine Translation*, pages 1049–1052, Online. Association for Computational Linguistics.
- Chau Tran, Shruti Bhosale, James Cross, Philipp Koehn, Sergey Edunov, and Angela Fan. 2021. [Facebook AI’s WMT21 news translation task submission.](#) In *Proceedings of the Sixth Conference on Machine Translation*, pages 205–215, Online. Association for Computational Linguistics.
- Arata Ugawa, Akihiro Tamura, Takashi Ninomiya, Hiroya Takamura, and Manabu Okumura. 2018. [Neural machine translation incorporating named entity.](#) In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3240–3250, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Shikhar Vashishth, Shyam Upadhyay, Gaurav Singh Tomar, and Manaal Faruqui. 2019. [Attention interpretability across nlp tasks.](#)
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need.](#) In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- David Vilar, Jia Xu, Luis Fernando D’Haro, and Hermann Ney. 2006. [Error analysis of statistical machine translation output.](#) In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC’06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Elena Voita, Rico Sennrich, and Ivan Titov. 2021. [Analyzing the source and target contributions to predictions 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 1126–1140, Online. Association for Computational Linguistics.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. [Context-aware neural machine translation learns anaphora resolution.](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. 2019. [Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned.](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Chaojun Wang and Rico Sennrich. 2020. [On exposure bias, hallucination and domain shift in neural machine translation.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3544–3552, Online. Association for Computational Linguistics.
- Sarah Wiegrefe and Yuval Pinter. 2019. [Attention is not not explanation.](#) 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 11–20, Hong Kong, China. Association for Computational Linguistics.
- Jianhao Yan, Fandong Meng, and Jie Zhou. 2022. [Probing causes of hallucinations in neural machine translations.](#)
- Chrysoula Zerva, Daan van Stigt, Ricardo Rei, Ana C Farinha, Pedro Ramos, José G. C. de Souza, Taisiya Glushkova, Miguel Vera, Fabio Kepler, and André F. T. Martins. 2021. [IST-unbabel 2021 submission for the quality estimation shared task.](#) In *Proceedings of the Sixth Conference on Machine Translation*, pages 961–972, Online. Association for Computational Linguistics.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert.](#) In *International Conference on Learning Representations*.
- Chunting Zhou, Graham Neubig, Jiatao Gu, Mona Diab, Francisco Guzmán, Luke Zettlemoyer, and Marjan Ghazvininejad. 2021. [Detecting hallucinated content in conditional neural sequence generation.](#) In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1393–1404, Online. Association for Computational Linguistics.

A Experimental setup

Data preprocessing. We filter our data using language identification and simple length-heuristics described in [Tran et al. \(2021\)](#). We encode the data with byte-pair encoding ([Sennrich et al., 2016](#)) using the SentencePiece framework ([Kudo and Richardson, 2018](#)). We set the vocabulary size to 32k and compute joint encodings and vocabulary.

Model parameters. We follow the setup of Transformer base model ([Vaswani et al., 2017](#)) (hidden size of 512, feedforward size of 2048, 6 encoder and 6 decoder layers, 8 attention heads). The model has approximately 77M parameters.

Optimizer. Similarly to ([Vashishth et al., 2019](#)), we train our model using the Adam optimizer with $\beta_1 = 0.9$ and $\beta_2 = 0.98$ and use an inverse square root learning rate scheduler with initial value 5×10^{-4} , and a linear warm-up in the first 4000 steps.

Training and Inference. Models are trained for 250K updates with a batch size of about 32K tokens. We set dropout to 0.3. At inference time, we produce translations using beam search with a beam of 5. We validate our models during training using SacreBLEU ([Post, 2018](#)), and we choose the checkpoint based on best BLEU in validation. We provide BLEU⁷ and COMET baselines on WMT evaluation campaigns in Table 3. We train and performance inference on top of the Fairseq framework ([Ott et al., 2019](#)).

COMET versions. We use models available in the official repository⁸: `wmt20-comet-da` for COMET and `wmt20-comet-qe-da-v2` for COMET-QE.

TokHal-Model. We follow the official implementation.⁹ For the synthetic data generation step, we used BART-large; and, for the token-level hallucination predictor, we used XLM-R ([Conneau et al., 2019](#)).

Computing Infrastructure. All our experiments have been ran on a machine with 2 physical Intel(R) Xeon(R) Gold 6348 @ 2.60GHz CPUs (total of 112 threads), and 4 NVIDIA RTX A6000 GPUs. In particular, the NMT model described above was trained in less than 30 hours on a single GPU.

⁷BLEU+case.mixed+lang.XX+numrefs.1+smooth.exp+tok.13a+version.1.4.2

⁸<https://github.com/Unbabel/COMET>

⁹<https://github.com/violet-zct/fairseq-detect-hallucination>

Metric	WMT2014	WMT2017	WMT2018
BLEU	31.1	32.6	38.9
COMET	0.3178	0.3257	0.3340

Table 3: Evaluation metrics for EN \rightarrow DE for *newstest* sets for the WMT 2014, 2017 and 2018 campaigns.

B Data Collection

We perform a rigorous and comprehensive manual annotation procedure with professional translators, in order to make sure that we are reliably analysing hallucinated translations.¹⁰ First, the translators were asked to be familiar with our task by reading our provided annotation examples, along with detailed annotation instructions. Then they had to pass a test to show they can recognize different pathologies in different translations. We selected the two best translators to gather the annotations. After being hired, we ran three one-hour tutorial sessions to explain the task thoroughly and to clarify any possible questions. During the annotation process, we made sure to be promptly available to answer any question from the translators. We paid a fair wage (25-30 USD per hour) – well above both the US federal minimum and the average EU minimum wage – and inspected their work for quality.

Guidelines. We make available the full guidelines used by the translators along with all other resources in the project repository. In short, the annotators were presented with a source sentence and a model-generated hypothesis and asked to deem that translation as correct (COR) or incorrect. If incorrect, they were prompted to answer a series of yes/no questions, regarding the presence of specific hallucinatory pathologies: oscillations (OSC), strong detachment (SD) and full detachment (FD). We also asked annotators to flag critical errors such as named-entity mistranslations (NE) and under-generated translations (UG).

Inter-annotator agreement. To determine the reliability of our annotations, we asked both our translators to annotate a set of 400 randomly sampled translations. For all hallucinatory categories but SD, the annotators achieved – according to Cohen’s kappa coefficient ([Cohen, 1960](#)) – almost

¹⁰Our translators were hired through Upwork and they were informed about the academic purposes of the data annotation process. All translators hired for this study reside in Europe.

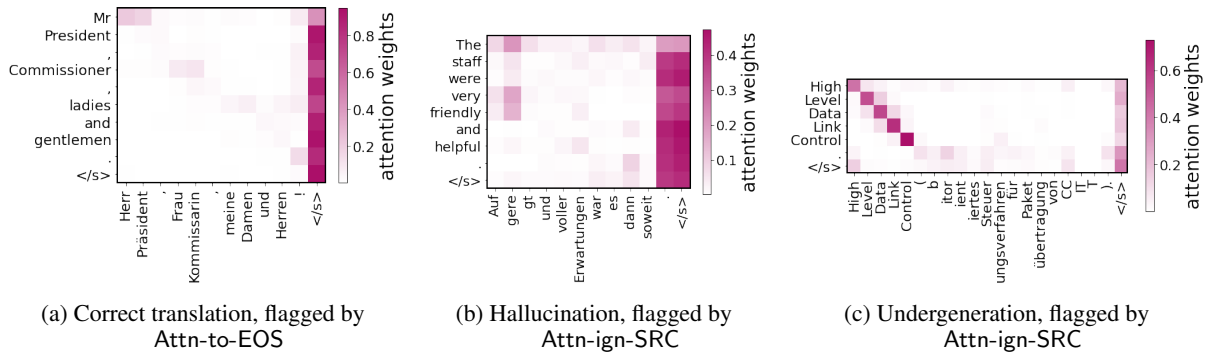


Figure 5: Examples of attention maps flagged by attention-based heuristics.

Fleiss's Kappa Scores					
COR	UG	NE	OSC	SD	FD
0.62	0.73	0.42	0.86	0.45	0.89

Table 4: Fleiss's Kappa inter-annotator agreement scores (\uparrow) for the different categories translators were prompted to identify.

perfect agreement. For all other categories, moderate to substantial agreement was obtained. This confirms that our data conforms very well to our instructions. The agreement scores for each category are displayed in Table 4.

C Patterns of attention maps for translations flagged with attention-based heuristics

The attention maps are shown in Figure 5. While both Attn-to-EOS and Attn-ign-SRC were designed to identify translations for which almost all the attention mass is concentrated on the EOS token, the patterns identified with Attn-ign-SRC are more diverse. They span from attention mass coming to various uninformative tokens (e.g., punctuation and other tokens as in Figure 5b) to examples shown in Figure 5c where attention is mostly diagonal (typically, these correspond to undergenerations; we follow this discussion in the next section).

D Analysing Less Severe Translation Errors

We notice that some of the detection methods are specialised on specific pathologies. For example, Attn-ign-SRC is by far the best for detecting undergenerations. This is expected: an undergeneration does not cover part of the source sentence, thus a significant proportion of source tokens receives little attention mass (see example in Fig-

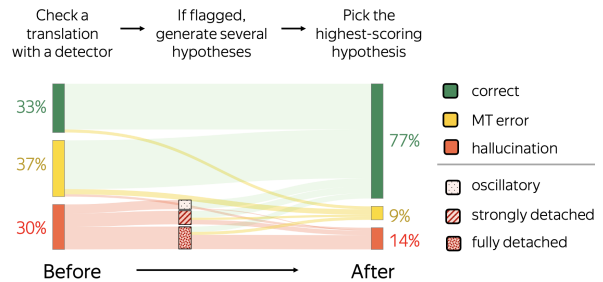


Figure 6: Our pipeline scheme along with results when we use Seq-Logprob as both the detector and the scoring measure.

ure 5c). For named entity errors, the best heuristic is TokHal-Model. This mirrors the discussion in Section 7.2: while severe errors (i.e., hallucinations) fall out of distribution for this model, mistranslations of short phrases are more in line with the model's training.

On a different note, MC-DSim performs much better for hallucinations than for less severe pathologies (Figure 3b). This again points to hallucinations being more unstable than other errors.

E Overwriting Hallucinations with Seq-Logprob as the scoring measure

The results are shown in Figure 6. In order to use Seq-Logprob as the scoring measure, we score all generated hypothesis with the original model. Overall, although the results follow the same trend as those using COMET-QE as the scoring measure, they are slightly worse: the hallucinatory rate is higher and the percentage of correct translations is smaller.

F Examples of overwritten translations

Table 5 shows examples of each hallucination type that have been overwritten with correct translations with the approach described in Section 9.

G Practical Recommendations for Detection

As we have mentioned in Section 3.4, Binary-label heuristics output a value in $\{0; 1\}$, and continuous-score heuristics output a value s in \mathbb{R} . These values s can be used to build a binary decision rule: for a given source x , translation \hat{y} and reference y , a translation is flagged by the detector if and only if $s(x, \hat{y}, y) \leq \gamma$. Naturally, s may not need be a function of x , \hat{y} and y (e.g. COMET-QE is only a function of x and \hat{y}).

Choosing the thresholds. In our work, we chose the thresholds γ for each detector by assessing the value s correspondent to approximately the 0.4-th percentile to be consistent with the data selection process (see Section 5). We computed these thresholds using the entire filtered held-out data. However, in practice, we obtained very similar results when we obtained these thresholds using a collection of only 10 000 examples from that dataset.

We recommend computing these thresholds on in-domain clean data. This will guarantee that the cut-off values were obtained in a scenario where the model performs best. Finally, the definition of the k -th percentile is expected to have an impact on the precision-recall trade-off. Thus, for preventive settings, we recommend sticking to more conservative values of k .

Choosing the detectors. The choice of detector rests upon the application for which it is intended. For high-precision settings, binary-label heuristics such as TNG and variants thereof (Raunak et al., 2022) may be more recommended. For preventive settings, we suggest using Seq-Logprob as the backbone of the hallucination detector. Naturally, we generally obtain higher recall by joining the sets of translations flagged with multiple methods. For example, Figure 3 reveals that joining the set of translations flagged with Seq-Logprob with the set of translations flagged with COMET-QE is very reasonable: COMET-QE performs better for oscillatory hallucinations, while Seq-Logprob is better for the other hallucination types.

Be careful when relying on references. Using reference information is helpful for detecting hallucinations (see Section 7.1), and while it may not be used to detect hallucinations on-the-fly, it may still prove useful for analysis works. We have found that high-quality parallel data is critical for adequate ap-

plication of these methods: very low scores might not only be attributed to poor translations, but also to reference mismatches. Indeed, preliminary experiments highlighted this worrying trend, which motivated us to clean the held-out set (Section 4). Thus, if using reference information to detect hallucinations, make sure to thoroughly clean your parallel data.

OVERWRITING FULLY DETACHED HALLUCINATIONS	
SOURCE REFERENCE	Handys, die bis auf Wäschewaschen und Staubsaugen scheinbar alles können. Mobile phones that can practically do everything except clean the laundry and vacuum clean.
ORIGINAL HYPOTHESIS	The staff were very friendly and helpful.
OVERWRITTEN HYPOTHESIS	Mobile phones that seem to be able to do everything except on laundry and dustproofing.
SOURCE REFERENCE	In unserem 2 Personen Van mit Dusche/WC war ausreichend Platz für uns beide. The space in our 2 person van with shower/toilet was enough for 2 people.
ORIGINAL HYPOTHESIS	The staff were very friendly and helpful. The room was clean and comfortable.
OVERWRITTEN HYPOTHESIS	In our 2 person van with shower/WC there was enough room for us both.
OVERWRITING STRONGLY DETACHED HALLUCINATIONS	
SOURCE REFERENCE	Tickets für Busse und die U-Bahn ist zu teuer, vor allem in Stockholm. Tickets for buses and the subway is too expensive, especially in Stockholm.
ORIGINAL HYPOTHESIS	The hotel is located in the centre of Stockholm, close to the train station.
OVERWRITTEN HYPOTHESIS	Buses and metro tickets are too expensive, especially in Stockholm.
SOURCE REFERENCE	Ich freue mich über jeden Tag, wo es mir so gut geht! I am pleased about each day, where I am so well!
ORIGINAL HYPOTHESIS	I look forward to seeing you every day!
OVERWRITTEN HYPOTHESIS	I'm very happy with every day I'm doing so well!
OVERWRITING OSCILLATORY HALLUCINATIONS	
SOURCE REFERENCE	In dieser Zeit stürzt sich Murnau bereits wieder in Theaterproben, allerdings widmet er sich nicht mehr der Schauspielerei, sondern der Regie. During this period, Murnau once again dedicates his time to theatre rehearsals; however, this time not as an actor, but as a director.
ORIGINAL HYPOTHESIS	Murnau was born in Murnau, Germany. He was born in Murnau, Germany.
OVERWRITTEN HYPOTHESIS	During this time, Murnau began to appear in theater tests again, but he was no longer concerned with acting, but with directing.
SOURCE REFERENCE	Die Staaten, deren Fangflotten im Nordwestatlantik Hochseefischerei betreiben, bemühen sich im Rahmen der NAFO (Organisation für die Fischerei im Nordwestatlantik) um eine gemeinsame Bestandserhaltung und -bewirtschaftung. States which fish in the high seas in the North West Atlantic co-operate in NAFO (North-west Atlantic Fisheries Organisation) in order to ensure conservation and management of stocks.
ORIGINAL HYPOTHESIS	The North-West Atlantic Fisheries Organisation (NAFO) is a member of the North-West Atlantic Fisheries Organisation (NAFO).
OVERWRITTEN HYPOTHESIS	The states whose fishing fleets in the North-West Atlantic are engaged in deep-sea fishing are seeking joint conservation and management within the framework of the North-West Atlantic Fisheries Organisation (NAFO).

Table 5: Examples of hallucinations of each type that have been overwritten with correct translations.