

Machine Translation Robustness to Natural Asemantic Variation

Jacob Bremerman Xiang Ren Jonathan May*

University of Southern California

Information Sciences Institute

{jbrem, xiangren, jonmay}@usc.edu

Abstract

Current Machine Translation (MT) models still struggle with more challenging input, such as noisy data and tail-end words and phrases. Several works have addressed this robustness issue by identifying specific categories of noise and variation then tuning models to perform better on them. An important yet under-studied category involves minor variations in nuance (non-typos) that preserve meaning w.r.t. the target language. We introduce and formalize this category as *Natural Asemantic Variation* (NAV) and investigate it in the context of MT robustness. We find that existing MT models fail when presented with NAV data, but we demonstrate strategies to improve performance on NAV by fine-tuning them with human-generated variations. We also show that NAV robustness can be transferred across languages and find that synthetic perturbations can achieve some but not all of the benefits of organic NAV data.

1 Introduction

While current Machine Translation (MT) models often perform decently in regular conditions, robustness to slight modifications on input remains a challenge. Several works have investigated the devastating effects that minor noise can cause in MT (Belinkov and Bisk, 2018; Khayrallah and Koehn, 2018). Addressing this is critical as robustness to naturally occurring variation is important in deployed systems.

Some classes of perturbations to which humans exhibit robustness but studies have shown MT models struggle with include spelling errors or rtypos, CaSinG (Niu et al., 2020), v15sua11y-51m1l14r ch4r4ct3r5 (Eger et al., 2019), and synonym replacement substitution (Wang et al., 2018). However, we focus on a less-known class we call Natural Asemantic Variation (NAV).

*Work done prior to JM joining Amazon.

Perturbation	Asemantic Variation
彼女は私に本を返した	informal verb conjugation
彼女は俺に本を返しました	masculine "me" pronoun
彼女は私に本を返してくれました	"favor" nuance

Table 1: Examples of NAV perturbations for a Japanese sentence, 彼女は私に本を返しました (she returned the book to me). These natural variations, which cause nuance differences in Japanese, cannot be translated succinctly into English. Therefore, all perturbations can be translated reasonably to the same English sentence. A robust MT model (or human) should be able to extract the same semantic meaning regardless of the variation provided.

Whereas aforementioned classes tend towards perturbations that generate ungrammatical, nonstandard, or unnatural sentences, NAV perturbations represent the extra-semantic linguistic properties of a language that allow for subtle changes in nuance while expressing the same core meaning. Specifically, a NAV perturbation is a kind of paraphrase in a source language that would not affect its translation in a given target language. Studying NAV is an important endeavor in MT robustness since it covers a less-contrived source of variation compared to existing MT robustness studies. As this topic is under-studied, it has also not been well-defined. We offer an example of NAV perturbations in Japanese in Table 1 and will provide further definitions and examples throughout the paper to help solidify the reader’s understanding.

From Table 1, we see how slight modifications of a sentence in one language can convey specific nuanced differences in that language, without warranting a change to its translation in another language. A common challenge in MT has to do with the fact that there are almost always several correct answers, and we can imagine translationally-equivalent mappings between large sets of similar sentences in different languages. Often we must resort to imperfect training and evaluating using subsets of these imagined sets due to normal data limitations. However, leveraging the STAPLE dataset (Mayhew

et al., 2020), we have a unique opportunity to investigate NAV phenomena more directly with access to several (often hundreds) of high-quality NAV perturbations. These perturbations allow us to test for NAV robustness, evaluating model behavior on challenging NAV test sentences.

In this paper, we contribute a formal definition for a linguistically-rich class of perturbations (NAV) and provide examples which help ground the concept (§2.1). We also contribute an evaluation setup to measure NAV robustness in the form of a repurposed test set from STAPLE and simple metrics to evaluate quality and consistency of MT models (§2.2). This formalization and metrics allow us to investigate specific research questions regarding NAV, as summarized below.

We evaluate existing MT models for NAV robustness and characterize the observed errors related to underspecification resolution and rarer expressions (§4.1). We show improvements in BLEU and our NAV-robustness metrics by exposing the model to NAV data and comparing various sub-selection and fine-tuning strategies (§4.2). We even find that these improvements can be transferred across languages, improving NAV robustness in languages that have not seen any NAV data (§4.3). We also design and perform synthetic perturbations, as these may provide a cheaper, more scalable way to improve NAV robustness without human annotations. We do achieve improvements though organic data is still more useful than synthetic (§4.4). Finally, we analyze our NAV-robust models in other evaluation settings with data that is out-of-domain and/or showing different classes of noise to investigate the side-effects of our methods (§4.5).

2 NAV Robustness

In robustness of NLP, it is difficult to determine improvements without considering both what types of input models are meant to be robust *to* and what robustness *looks like*. In this section, we will first discuss NAV with a formal definition and explanatory examples. Then we will define NAV robustness in terms of desired behaviors and present metrics as proxies for measuring NAV robustness.

2.1 NAV Perturbations

To formally define *natural asemantic variation*, we consider a corpus C in two languages \mathcal{X}, \mathcal{Y} that contain c paired clusters of translationally-

1	Je vais à la bibliothèque → I'm going to the library Il va à la bibliothèque → He's going to the library	Non-NAV _{fr,en} : encodable semantic shift
2	Je vais à la bibliothèque → I'm going to the library Je vais à la Bibliothèque → I'm going to the library	Non-NAV _{fr,en} : unnatural, agrammatical variant
3	Je vais à la bibliothèque → I'm going to the library Moi je vais à la bibliothèque → I'm going to the library	NAV _{fr,en} : same English translation
4	Je vais à la bibliothèque → Voy a la biblioteca Moi je vais à la bibliothèque → Yo voy a la biblioteca	Non-NAV _{fr,es} : different Spanish translation

Figure 1: 4 examples of perturbations to an original French sentence ‘Je vais à la bibliothèque’. Only one perturbation (example 3) can be called a NAV perturbation as the variation is an emphasis on the subject, difficult to encode in written English. The same perturbation is however not NAV in the context of a different target language that does easily encode subject emphasis (example 4).

equivalent natural and grammatical sentences:

$$C = \{(X_i, Y_i)\} \quad \forall i \in \{1, \dots, c\};$$

$$X_i = \{x_i^1, \dots, x_i^n\}; \quad Y_i = \{y_i^1, \dots, y_i^m\};$$

$$\forall j \in \{1, \dots, n\}, \forall k \in \{1, \dots, m\}, x_i^j \Leftrightarrow y_i^k,$$

where \Leftrightarrow represents *translational equivalence*, meaning the two sentences x_i^j and y_i^k , which are in clusters X_i and Y_i , respectively, can be reasonably translated to each other. For example, all the Japanese examples in Table 1 are translationally equivalent with *she returned the book to me* in English. Another valid y could be *she gave the book back to me*. All $x \in X_i$ can then be considered NAV _{\mathcal{X}, \mathcal{Y}} perturbations of each other. A NAV perturbation must be considered in context of the language pair. x_i^1 and x_i^2 may be NAV _{\mathcal{X}, \mathcal{Y}} perturbations of each other but not NAV _{\mathcal{X}, \mathcal{Z}} perturbations.¹

In this work, we further limit the corpus to be:

$$\forall i \quad |X_i| > 1, \quad |Y_i| = 1,$$

meaning only one reference translation is provided. This choice is due to data limitations as we only have access to semantically-equivalent variations on the source side. (if $|X_i| = 1$ as well, then C has the form of most standard MT datasets).

To ground these definitions, Figure 1 offers specific examples to distinguish NAV perturbations

¹Note: A semantic paraphrase x_i^1 of a sentence x_i^2 in the same language \mathcal{X} cannot be assumed to be a NAV perturbation *on its own* because it may or may not cause a change to its translation in the relevant target language.

from non-NAV perturbations. The modifications in examples 1 and 2 are either nonstandard or have too large of an impact on semantics such that their translations are different. Example 3 shows a simple example of a NAV perturbation while example 4 demonstrates the language dependence of NAV by using the same perturbation from 3, highlighting the complexity of this problem. For a different target language, the two sentences would no longer exist in the same X_i .

2.2 NAV Robustness

In order to evaluate models for their robustness to NAV perturbations, we consider two desiderata, *quality* and *consistency*, which we evaluate separately. There are often several high-quality answer options in MT, so two models could both produce high-quality translations but differ greatly in their consistency (how similar their hypotheses among minimally-altered, translationally-equivalent inputs are). For our purposes, we posit that a more NAV-robust MT model should not only have increased or maintained translation quality but also consistency in output when given different NAV perturbations as input. Consistency in output would help ensure overall system robustness if we imagine the MT system being part of a larger NLP pipeline.

In this work, for translation quality, we use BLEU (Papineni et al., 2002) with SacreBLEU² (Post, 2018). For translation consistency we define a metric, $CONSIST^\theta$, which rewards a translation model f_θ for exhibiting less variation in its outputs $\hat{Y}_i^\theta = f_\theta(X_i)$ among NAV-perturbed inputs X_i . Specifically, for a pair (X_i, \hat{Y}_i^θ) , $CONSIST^\theta$ is calculated as follows:

$$CONSIST_i^\theta \triangleq \frac{1}{|X_i|} \sum_{j=1}^{|\hat{Y}_i^\theta|} \frac{|\hat{y}_i^j|}{j},$$

where $\hat{Y}_i^\theta \triangleq \{\hat{y}_i : \exists x_i^j \in X_i \wedge \theta(x_i^j) = \hat{y}_i\}$, $|\hat{y}_i| \triangleq |\{x_i : x_i \in X_i \wedge \theta(x_i) = \hat{y}_i\}|$, and we order \hat{Y}_i^θ as $[\hat{y}_i^1, \dots, \hat{y}_i^m]$, sorted descending by $|\hat{y}_i^j|$. Intuitively, $|\hat{y}_i^j|$ is the number of $x_i \in X_i$ for which $\theta(x_i)$ outputs hypothesis \hat{y}_i^j . We then average this score across all translation cluster-pairs in the corpus:

$$CONSIST^\theta \triangleq \frac{1}{c} \sum_{i=1}^c CONSIST_i^\theta.$$

²`nrefs:1|case:lc|eff:no|tok:13a|smooth:exp|version:2.0.0`

We drop the superscript when unambiguous. This metric may break down in long-output tasks but works well for our domain of short, simple sentences. Essentially, this metric works by punishing the model more and more for novel translation outputs on translationally-equivalent inputs. An explanatory example calculation is available in Appendix §A.1. We also discuss an alternative consistency metric with trade-offs in Appendix §A.2.

3 Analysis Setup

In this work, we seek to provide answers to several analysis questions related to NAV perturbations and MT models’ robustness to them. Using our quality and consistency metrics we investigate the following:

- Are existing MT models robust to NAV? if not, what types of errors are the models making?
- How can NAV perturbation data be used to improve an MT model’s NAV robustness?
- Can a model’s improved NAV robustness in one language pair be transferred to another language pair?
- How do synthetic perturbations compare to organic NAV data w.r.t. NAV robustness?
- What behavior changes does a ‘NAV-robustified’ MT model exhibit in other MT contexts such as out of domain data or on other types of robustness tests?

3.1 Baseline Models

We work with two main classes of baseline models: *mono-pair* (can translate one specific language to one specific other language, e.g. ja-en) and *multi-pair* translation models (can translate between several languages, e.g. {hu,ja,pt}-{hu,ja,pt}). All models are transformer models (Vaswani et al., 2017) using the fairseq toolkit (Ott et al., 2019). Our principal models are multi-pair since many of our analysis questions relate to language transfer, but we also perform preliminary experiments with mono-pair models for completeness.

We pre-train mono-pair models with the Tatoeba corpus for {hu,ja,pt}-en. This corpus is fairly close to the STAPLE domain and allows for reasonable baseline models. We obtain subwords and tokenize using SentencePiece (Kudo and Richardson, 2018).

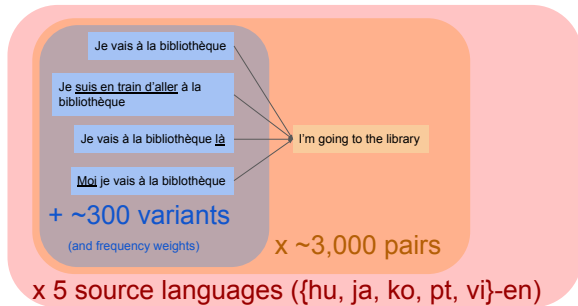


Figure 2: Diagram describing the STAPLE data. French is used to mirror the French examples used earlier for ease of relating NAV concepts.

Hyperparameters are the default settings for the `transformer` architecture in `fairseq`.

For multi-pair model experiments, we use the M2M-100 model (Fan et al., 2021), which is trained on CC-Matrix (Schwenk et al., 2021) and CC-Aligned (El-Kishky et al., 2020). This model comes with a predetermined vocabulary and tokenizer. It uses the `transformer_wmt_en_de_big` architecture in `fairseq` with 24 encoder and 24 decoder layers. The checkpoints available online are used as initialization for further fine-tuning using different NAV data, as described next.³

3.2 NAV Data

This research is made possible largely by the unique dataset publicly released by Duolingo for the STAPLE shared task (Mayhew et al., 2020). For each English sentence, several (average 300) accepted translations of that sentence in one of five languages were sourced from Duolingo users and further annotated with frequency scores based on how often they used that specific translation. We use the enumerated valid translations as source-side NAV perturbations translating to the same target English sentence. A diagram representing the data can be seen in Figure 2. Table 1 includes a real example from the training split for Japanese-English.

For NAV fine-tuning (§4.1), we use different subsets of the STAPLE training corpus in three language pairs {hu,ja,pt}-en. For each “many-to-one” pair, we either use **all** of the many NAV perturbations or **subselect**. When **subselecting**, we consider three strategies. Since we also have frequency scores for each perturbation, we could select a number of the **most** frequent perturbations, **least** frequent perturbations, or uniformly sample

³More details on all the specifications for our models are available in experiment scripts released publicly in our [github](#)

lang	train	dev	test
hu	250k/4k	28k/500	34k/500
ja	860k/2.5k	170k/500	170k/500
pt	530k/4k	60k/500	68k/500
ko	-/-	-/-	150k/500
vi	-/-	-/-	28k/500

Table 2: Statistics for STAPLE data used in our experiments. *unique source segments / unique target segments*

for a **random** subset. In our paper, a subset is defined by a number and a strategy, so “10-random” means 10 random NAV perturbations were chosen for each example.

For NAV robustness evaluation, we use the full (all NAV perturbations included) many-to-1 pairs of the STAPLE test split. On top of the {hu,ja,pt}-en sets, we hold out {ko,vi}-en test sets to evaluate 0-shot language transfer using multi-pair models (§4.2). Data statistics are shown in Table 2.

3.3 Synthetic Data Augmentation

Obtaining NAV perturbations is expensive, requiring human annotation by bilingual language experts to generate natural variations that don’t affect translation equivalence. We consider applying synthetic perturbations to some of the STAPLE data to compare potential NAV robustness gains with a cheaper, more scalable strategy.

With the organic, human-generated variations in the STAPLE data, we create a synthetic dataset by taking the 1-most frequent source sentence from each translation pair and perturbing it synthetically 9 times for a total of 10 variations (while the 10 copies of the target sentence are left unchanged).

For our roman script languages {hu,pt}-en, we use known noising techniques of random casing changes and character substitutions, insertions and deletions as in Niu et al. (2020). These perturbations are not NAV perturbations, but they represent a common and easy way to add noisy data.

We also add synthetic NAV perturbations for the ja-en split. Roman script noising strategies are not perfectly translatable to ja-en. Also, ja-en arguably has the easiest rule-based modifications that can be programmed to automatically generate NAV perturbations. By no means are these methods exhaustive, but a combination of simple insertions of emphasis-related particles, substitutions of pronouns based on gender identity, and dropping unnecessary pronouns serve as a basic technique for synthetic NAV noising. Rules were implemented after scanning

training data and observing patterns of NAV perturbation.⁴

3.4 Additional Evaluations

We use other existing MT datasets to evaluate our models in contexts such as performance on out-of-domain (OOD) data (§B.1) and robustness-transfer to other types of noisy input (§B.2). We use test splits from Tatoeba, OPUS-100 (Zhang et al., 2020) and MTNT (Michel and Neubig, 2018).

We also create additional evaluation sets from the STAPLE data to analyze in-domain model behavior. In addition to the ‘all perturbations’ test sets, we also filter out ‘1-most’ and ‘1-least’ test sets. These provide a more standard MT evaluation framework that relies on BLEU to measure in-domain translation performance on common (‘1-most’) and robustness to uncommon, NAV-perturbed (‘1-least’) sentence pairs.

4 Experiments and Results

For our experiments and results section, we will progress chronologically as we address each of our analysis questions as defined in Section 3.

4.1 Existing MT NAV Behavior

Before we perform comprehensive experiments to analyze NAV robustness, we investigate how existing MT models perform on NAV perturbations. This way we can establish 1) if this even is a problem in MT and 2) if so, what types of errors are we aiming to address? We simply take our baseline models (§3.1) and search for error patterns produced from example inputs in the STAPLE dataset.

One pattern of errors involves an inability to properly resolve cases of under-specification. For example, the M2M-100 model properly handles “それはトマトの種類です。”, outputting “This is a type of tomato.” However, a NAV-perturbed input, “トマトの種類です。” causes the model to output ‘Species of tomatoes.’ since this sentence drops the subject, common in Japanese.

Another pattern of errors involves an inability to handle less common synonyms or paraphrases for the same target word(s). ‘Please’ is the main way to express politeness in a request in English but other languages have multiple common ways to express this. The M2M-model properly translates the Portuguese sentence “por favor, não a desperte

⁴Perturbation scripts are available on [github](#). Rules are described in Appendix §A.3.

agora.” as “Please don’t wake her up now.” but mistakenly translates “por **obséquoio**, não a desperte agora.” as “For obsession, don’t wake her up now.”.

In just a small sample of examples, we find several instances of these types of errors, confirming our intuition that existing MT models cannot handle NAV perturbations completely. We are also able to see that these types of mistakes will be captured by our BLEU and CONSIST metrics, motivating our quantitative experiments. More information on this error analysis is available in Appendix §A.6.

4.2 Improving NAV Robustness with NAV Data

For our first experiments, we fine-tune our mono-pair transformer models pre-trained on Tatoeba using strategies for subselecting STAPLE data discussed in Section 3.2 and evaluate using metrics from Section 2.2. After initially experimenting with several combinations of number of perturbations per set pair and selection strategy,⁵ we report on 4 conditions, representative of common “real-world” MT scenarios.⁶

baseline: off-the-shelf MT model, no fine-tuning (§3.1)

+ 1-most: simple domain adaptation using parallel text in target domain (one “typical” translation per sentence).

+ 10-random: preferred “NAV robustness” data condition, having a small yet diverse set of NAV perturbations per translation pair

+ all: “throw in all the data” approach

Our results from these experiments are shown in Figure 3. The main purpose of these preliminary experiments is to provide evidence that NAV robustness improvement is possible on smaller scale models. We test various configurations to justify design choices for our future, more comprehensive experiments with large multi-pair models. We find that the ‘+ 10-random’ condition results in the best trade-off of BLEU and CONSIST.

We repeat these experiments using a large multi-pair model. We find similar patterns in the results, confirming “10-random” as a decent strategy for improving NAV robustness. Here, the baseline is the M2M-100 model and models are fine-tuned

⁵full results in appendix

⁶We choose to only report on these conditions to strike a balance between not overloading the reader with too many experiment settings and providing enough that useful conclusions can be drawn by comparing.

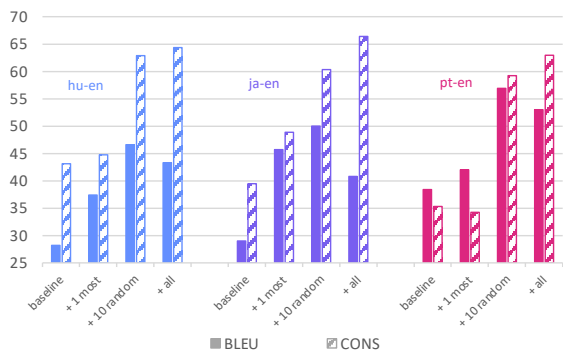


Figure 3: Results from mono-pair model experiments. The ‘+ 1 most’ condition shows improvements over a baseline in a typical MT domain adaptation scenario where some target domain data is available. The ‘+ 10 random’ condition shows how exposing the model to NAV perturbations further increases both BLEU and CONSIST scores beyond domain adaptation with frequent input forms. While including all NAV perturbations (‘+ all’) continues to improve CONSIST, we see a decrease in BLEU.

and evaluated on only the same language pair.⁷ As the patterns are similar, results are plotted in the Appendix A.5.

4.3 Transferring NAV Robustness across Languages

With a large multi-pair model, we can consider zero-shot language-transfer of NAV robustness. First, we take the M2M-100 baseline models and fine-tune them on one of the STAPLE fine-tuning language-pairs. We then evaluate on our held-out language test sets (unseen pairs during fine-tuning, {ko,vi}-en). Our results are shown in Figure 4. From previous experiments, we see sufficient evidence that “+ 10 random” is appropriate for NAV robustifying, so we mainly report on this setting.

Results suggest zero-shot transfer of NAV robustness is possible, with 10-random NAV perturbations per pair showing larger robustness improvement than simply fine-tuning on 1-most. We also observe language differences. ko-en BLEU improves more from ja-en fine-tuning while vi-en improves more from hu-en.

Finally, we perform multilingual fine-tuning experiments in which we combine all three training sets together {hu,ja,pt}-en with the ‘+ 10 random’ strategy and evaluate on all 5 test sets. We compare

⁷To help validate our results, we run 3 different seeds for the randomization of data in ‘random-10.’ Standard deviation for BLEU scores is between 0.4 and 0.6. Standard deviation for CONSIST scores is between 0.7 and 1. The variation is small enough to not affect conclusions.

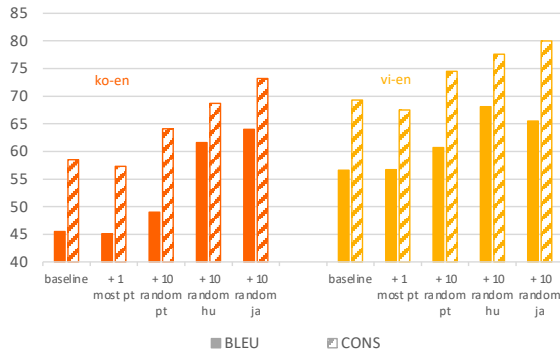


Figure 4: As in mono-pair experiments, including 10 NAV perturbations per pair shows larger robustness improvement than simply fine-tuning on 1-most.

the multilingual fine-tuning results to the previous best results (according to BLEU) from our monolingual fine-tuning experiments. Results are shown in Figure 5. We see that using all 3 test sets for NAV fine-tuning is consistently a better option over any given single set.

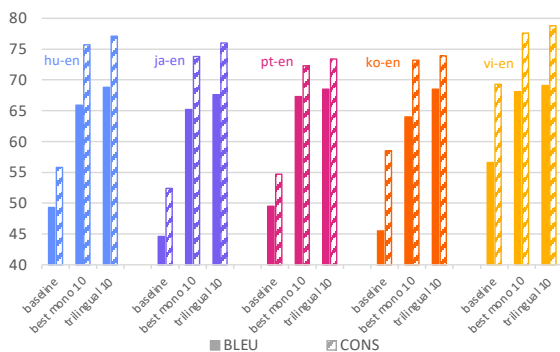


Figure 5: Multilingual (trilingual) fine-tuning improves both BLEU and CONSIST for every evaluation language pair compared to that pair’s previous best model using monolingual fine-tuning.

4.4 Approaching NAV Robustness with Synthetic Data Augmentation

We repeat experiments from the previous section using our synthetic fine-tuning sets (Section 3.3) and compare to the organic 10-random sets. We also combine all the synthetic sets {hu,ja,pt}-en for multilingual fine-tuning and evaluate on our NAV robustness evaluation sets. Results are shown in Figure 6.

Our synthetic fine-tuning sets do improve robustness compared to a baseline model, but they are far from achieving the same improvements as the organically generated data. We also see how our NAV-oriented synthetic data (ja-en) more closely

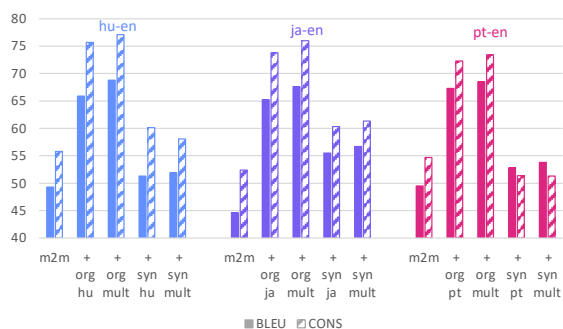


Figure 6: BLEU and CONSIST scores for organic vs. synthetic perturbations during fine-tuning. Synthetic perturbations improve robustness over baselines but organic perturbations are much more useful.

approximates organic NAV data gains compared to our non-NAV synthetic data ($\{hu,pt\}$ -en).

We also test synthetic fine-tuning performance in zero-shot language transfer. Results are shown in Figure 7. We continue to see that synthetic data is not able to provide the same NAV robustness gains as organically-generated data. However, in zero-shot language transfer, multilingual synthetic data more closely approaches organic-data performance than it can with in-language fine-tuning. This suggests in-language organic NAV data is most useful even if some NAV robustness can be improved with transfer.

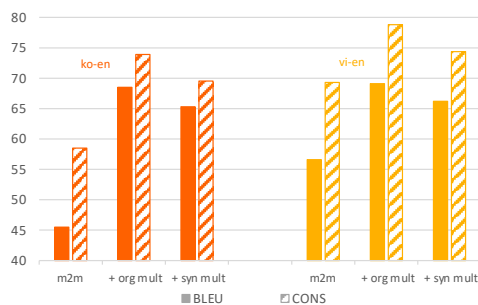


Figure 7: BLEU and CONSIST scores for organic vs. synthetic perturbations during fine-tuning in zero-shot language transfer. Synthetic perturbations improve robustness over baselines but organic perturbations are much more useful. However, this gap is smaller than that shown with in-language experiments.

4.5 Additional Evaluations

Results from the previous sections suggest that NAV robustness can be effectively learned (even in zero-shot scenarios) by exposing models to a variety of several NAV perturbations per translation example. However, these findings raise new questions about what specifically the models are

learning, how transferable that learning is and how else the model’s behavior changes.

We perform several analysis experiments in which we take our best models from previous experiments, which we designate as our ‘NAV-robust’ models and evaluate them compared to baseline MT models in other evaluation settings. These new test settings include OOD data (§B.1), other classes of noise (§B.2) and more frequent (less NAV-relevant) in-domain data (§B.3).

The overall takeaway from these experiments is that our NAV-robust models sacrifice some performance in non-NAV-relevant settings. This is not uncommon in general NLP fine-tuning, as often models fine-tuned for a new task, domain, etc. often underperforms in the original setting (Thompson et al., 2019; He et al., 2021).

These experiments serve primarily to confirm that there are limitations and drawbacks to improving NAV robustness with our current proposed methods. We include specifics in the Appendix to distill the main work for novel contributions. These results suggest avenues for future work to avoid these trade-offs for robustness, which remains a problem in the wider NLP community.

One set of results we communicate here is performance on in-domain ‘1-most’ and ‘1-least’ test data. These test settings are useful because they allow us to directly compare models on similar text examples that differ in frequency of use, thus distinguishing between ‘normal’ sentences and heavily-NAV-perturbed sentences. The latter evaluation serves as another proxy for robustness.

In most of our NAV experiments from earlier sections, ‘+ 10 random’ consistently performs better than ‘+ 1 most’. This is reflected in the ‘1-least’ evaluation as well. However, in the ‘1-most’ evaluation, ‘+ 1 most’ outperforms ‘+ 10 random’. What this shows is that evaluating on frequent, non-noisy data may select for models that actually perform worse on data of the same domain but with less frequent NAV forms. These results can be seen for ja-en in Figure 8 and full results (which mirror those shown here) are shown in §B.3.

5 Discussion

Overall our work raises several questions about NAV phenomena and how to address them in MT. We are able to provide some answers to a few research questions and open pathways for future work. NAV describes a subtle yet substantial mark

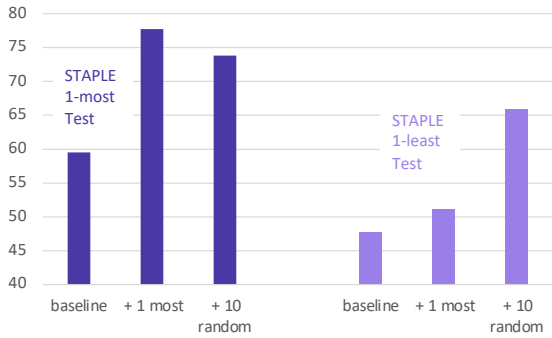


Figure 8: BLEU scores evaluating on 1-most and 1-least STAPLE test sets in ja-en. Fine-tuning with common ‘1-most’ data is better than using NAV perturbations for testing on 1-most inputs. However, testing on rarer inputs exhibiting NAV phenomena in the 1-least test data, intentional NAV robustifying (fine-tuning with ‘+ 10 random’ data) is necessary over the common 1-most data.

	hu		ja		pt		ko		vi	
Tatoeba Transformer	28.2	43.1	29.0	39.5	38.4	35.31	-	-	-	-
+ 1 most	37.4	44.8	45.7	48.9	42.0	34.24	-	-	-	-
+ 10 random	46.6	62.9	50.0	60.3	56.9	59.22	-	-	-	-
+ all	43.3	64.4	40.8	66.4	53.0	62.95	-	-	-	-
M2M-100	49.3	55.8	44.6	52.4	49.5	54.7	45.5	58.5	56.6	69.3
+ 1 most	52.7	60.6	55.8	60.0	51.2	47.3	59.4	66.1	64.5	71.8
+ 10 random	65.9	75.7	65.2	73.8	67.3	72.3	64.0	73.2	68.1	77.6
+ all	65.5	78.8	55.3	75.2	66.7	77.6	59.0	74.8	66.0	78.8
+ multi organic	68.8	77.1	67.6	76.0	68.5	73.4	68.5	73.9	69.1	78.8
+ synthetic	51.3	60.1	55.5	60.3	52.8	51.4	60.1	65.8	64.5	74.8
+ multi synthetic	51.9	58.1	56.7	61.3	53.8	51.3	65.3	69.5	66.2	74.4

Table 3: Main results from all NAV experiments organized by source language (columns) and baseline/fine-tuning method (rows). In each column, the left blue elements are BLEU scores, the right black elements are CONSIST scores.

of language that stands as an appropriate target for current NLP research.

Addressing NAV in a systematic way is an incredible challenge due to the vastness of possible variation. The STAPLE dataset provides a suitable way to begin investigating solutions to NAV-related problems in MT. We are able to show how NAV robustness can be improved using STAPLE data and how those improvements can transfer across languages. There seems to be language dependence whereby robustness learned in one language can have varying effects depending on the language transferred to.

We find that NAV-robust models can perform worse in non-NAV settings, which does not differ from several other findings in robustness work in which a more robust model may not perform better on original, non-noised input. We also address scaling issues by synthetically generating NAV and non-NAV examples. The synthetic NAV examples seem to help more than the non-NAV examples, but

none of them are able to obtain the same improvements as human-created NAV data.

The major results from all of our experiments are displayed in Table 3 for side-by-side comparison.⁸

6 Related Work

Paraphrasing Paraphrasing has been shown to be beneficial in MT contexts. Hu et al. (2019) have done work in releasing and using paraphrase data. Khayrallah et al. (2020) use generated paraphrases as data augmentation for MT. One thing that differentiates our work is the formal definition of NAV. It can be difficult to clarify what is or is not a paraphrase, and this work attempts to formalize a specific kind of paraphrase.

MT Robustness Several works investigate robustness in MT by considering different classes of perturbations and developing strategies to improve performance at test time. Niu et al. (2020) perform misspelling- and casing-related perturbations on MT test input and evaluate different models on their robustness to these. Salesky et al. (2021) address several classes of perturbations by replacing a standard unicode-based encoder with a visual encoder, which demonstrates higher robustness to perturbations associated with visual appearance of language, such as 1337speak. Zhang et al. (2020) generate additional clean and noisy data using back-translation and different datasets to improve performance on noisy text. We add to this line of work by addressing a specific kind of robustness (NAV).

Data Augmentation There has also been substantial work in using synthetic data to improve MT more generally. Sennrich et al. (2016) use back translation to augment MT datasets, which improves performance. Wang et al. (2018) augment data by reusing sentences with synonyms replaced. Karpukhin et al. (2019) and Berard et al. (2019) improve robustness by adding synthetic noise. Our work is unique in our attempts to augment data using developed rules based on asemantic under-specification in language.

7 Conclusion

We present the reader with an under-studied subject in MT and contribute more formalized definitions

⁸This table offers a concise, overarching view of this work’s experiments, as contrasted with the extracted figures in Section 4, which serve to highlight evidence pertaining to specific research questions.

and simple evaluation metrics for the phenomenon of natural asemantic variation. We also perform experiments, showing how NAV perturbations can be used during fine-tuning to improve robustness of MT models, even when evaluating on a different language-pair. Several questions remain to be explored to further formalize NAV, investigate its role in NLP modeling and improve techniques to increase nuanced understanding.

8 Limitations

Several of the limitations of this work are discussed throughout the paper and in the Appendix. To summarize, the metrics used to evaluate NAV robustness may not work as well in other domains such as data with longer sentences. However, the introduced formalism of NAV is relevant to any domain.

A limitation to mention about the formalism is that it likely does not perfectly distinguish NAV and non-NAV sentences. Just as translators can argue about translation decisions, there are arguments for why or why not certain sentences should be considered ‘translationally-equivalent’. ‘Asemantic’ is also a term that has gray areas. Nonetheless, the formalism allows us to establish metrics and a foundation for analysis for this very real phenomenon.

Our methods are also limited since they rely on access to unique high-quality data. We address this with language-transfer and synthetic perturbations, but we are unable to obtain the same benefits as human-generated in-language data.

Other limitations include language coverage. We work with 6 languages, which offers more support for generalization than only one language pair but still could raise some doubts about applicability to all languages.

There are also equipment limitations to consider. Most of our experiments use the M2M-100 model as a baseline, which required our experiments to be run on RTX 6000 GPUs. We do run some of the experiments on smaller baseline models which fit on older RTX 2800 machines and see similar results.

Acknowledgments

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Agreement No. HR00112290056. Xiang Ren’s research is supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research

Projects Activity (IARPA), via Contract No. 2019-19051600007 and NSF IIS 2048211.

References

- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *International Conference on Learning Representations*.
- Alexandre Berard, Ioan Calapodescu, Marc Dymetman, Claude Roux, Jean-Luc Meunier, and Vassilina Nikoulina. 2019. [Machine translation of restaurant reviews: New corpus for domain adaptation and robustness](#). In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 168–176, Hong Kong. Association for Computational Linguistics.
- Steffen Eger, Gözde Gül Şahin, Andreas Rücklé, Ji-Ung Lee, Claudia Schulz, Mohsen Mesgar, Krishnkant Swarnkar, Edwin Simpson, and Iryna Gurevych. 2019. [Text processing like humans do: Visually attacking and shielding NLP systems](#). 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 1634–1647, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. 2020. [CCAligned: A massive collection of cross-lingual web-document pairs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5960–5969, Online. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. [Beyond english-centric multilingual machine translation](#). *Journal of Machine Learning Research*, 22(107):1–48.
- 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.
- Tianxing He, Jun Liu, Kyunghyun Cho, Myle Ott, Bing Liu, James Glass, and Fuchun Peng. 2021. [Analyzing the forgetting problem in pretrain-finetuning of open-domain dialogue response models](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1121–1133, Online. Association for Computational Linguistics.

- J. Edward Hu, Rachel Rudinger, Matt Post, and Benjamin Van Durme. 2019. [Parabank: Monolingual bitext generation and sentential paraphrasing via lexically-constrained neural machine translation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6521–6528.
- Vladimir Karpukhin, Omer Levy, Jacob Eisenstein, and Marjan Ghazvininejad. 2019. [Training on synthetic noise improves robustness to natural noise in machine translation](#). In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 42–47, Hong Kong, China. Association for Computational Linguistics.
- Huda Khayrallah and Philipp Koehn. 2018. [On the impact of various types of noise on neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 74–83, Melbourne, Australia. Association for Computational Linguistics.
- Huda Khayrallah, Brian Thompson, Matt Post, and Philipp Koehn. 2020. [Simulated multiple reference training improves low-resource machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 82–89, Online. Association for Computational Linguistics.
- 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.
- Stephen Mayhew, Klinton Bicknell, Chris Brust, Bill McDowell, Will Monroe, and Burr Settles. 2020. [Simultaneous translation and paraphrase for language education](#). In *Proceedings of the Fourth Workshop on Neural Generation and Translation*, pages 232–243, Online. Association for Computational Linguistics.
- Paul Michel and Graham Neubig. 2018. [MTNT: A testbed for machine translation of noisy text](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 543–553, Brussels, Belgium. Association for Computational Linguistics.
- Xing Niu, Prashant Mathur, Georgiana Dinu, and Yaser Al-Onaizan. 2020. [Evaluating robustness to input perturbations for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8538–8544, 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.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. 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.
- Elizabeth Salesky, David Etter, and Matt Post. 2021. [Robust open-vocabulary translation from visual text representations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7235–7252, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, and Angela Fan. 2021. [CCMatrix: Mining billions of high-quality parallel sentences on the web](#). 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 6490–6500, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. [Overcoming catastrophic forgetting during domain adaptation of neural machine translation](#). 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 2062–2068, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. [SwitchOut: an efficient data augmentation algorithm for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 856–861, Brussels, Belgium. Association for Computational Linguistics.

Biao Zhang, Philip Williams, Ivan Titov, and Rico Sen-
nrich. 2020. [Improving massively multilingual neu-
ral machine translation and zero-shot translation](#). In
*Proceedings of the 58th Annual Meeting of the Asso-
ciation for Computational Linguistics*, pages 1628–
1639, Online. Association for Computational Linguis-
tics.

A Appendix

A.1 CONSIST Metric Example

Say, as an example, we have 1000 NAV perturbations of the same source sentence. The model will then have 1000 output sentences, many of which will (hopefully) be the same. We can then create groups of identical translation outputs and rank them by size from largest to smallest. In our example, the model outputs 3 different translations, one 750 times, one 200 times, and one 50 times. We define *CONSIST* as:

$$CONSIST = \frac{1}{n} \sum_{i=1}^n \frac{|G|_i}{i}$$

where i is the rank of the group, n is the number of total input-output example pairs and $|G|_i$ is the ratio of the size of the i th largest group to n . In our example, $CONSIST = \frac{1}{1000}(\frac{750}{1} + \frac{200}{2} + \frac{50}{3})$. In this way, models are rewarded for producing a small number of groups with large sizes, suggesting higher consistency.

A.2 Pair-Wise BLEU

The CONSIST metric we propose has several strengths. It ensures very fine-grained similarity, which would benefit NLP systems in which MT is part of a larger pipeline. In such a case, tiny variations could cause large down-stream errors. CONSIST is also quick to compute since it does not require comparing each hypothesis with each other one separately. However, we concede that CONSIST may not be best in all settings, especially for domains with longer sentences.

To address this weakness, we also calculate a pairwise-BLEU-based score in which we compute average sentence BLEU between all pairs of hypotheses for a prompt and average across all prompts. Specifically, for a pair (X_i, \hat{Y}_i^θ) :

$$PWB_i^\theta \triangleq \frac{1}{n_i} \sum_{j=1}^{|\hat{Y}_i^\theta|} \sum_{k=j+1}^{|\hat{Y}_i^\theta|} sBLEU(\hat{y}_i^j, \hat{y}_i^k),$$

$$n_i = \frac{(|X_i|)(|X_i| - 1)}{2}$$

Thus a model’s final PWB score is:

$$PWB^\theta \triangleq \frac{1}{|C|} \sum_{i=1}^{|C|} PWB_i^\theta.$$

	CONSIST	PWB
m2m	54.7	58.3
+ 1 most	47.3	53
+ 10 random	72.3	73.2
+ all	77.6	78.3

Table 4: CONSIST scores match patterns with the more robust yet more expensive PWB metric for consistency in pt-en.

This metric is like Minimum Bayes Risk (Freitag et al., 2022) and like that metric, takes significantly longer than CONSIST to compute. It mitigates some of the negatives of CONSIST since it works for long sentences and allows for less-strict similarity evaluation.

We observe the same consistency patterns from our models using this pairwise score as the patterns drawn from CONSIST shown in Table 4.

The similarity in consistency scores between these two metrics motivates our usage for the quicker of the two, CONSIST.

A.3 Synthetic Perturbation

A.3.1 Latin Script Perturbations (hu, pt)

For the Latin-script perturbations, each word (determined by white space) has a uniform random-chance of one of the following perturbations:

- No change
- Casing e.g. ‘apple’ → ‘Apple’
- Insertion e.g. ‘apple’ → ‘appqle’
- Deletion e.g. ‘apple’ → ‘aple’
- Substitution e.g. ‘apple’ → ‘apqle’

A.3.2 Japanese Script NAV Perturbation

For the Japanese-script perturbations, we use rules (which also have uniform chance of being implemented after certain conditions are met). These rules attempt to create true NAV perturbations, compared to the more-contrived Latin-script perturbations. Rules include:

- Sentence-Ending Emphasis Particles e.g. “ → ‘よ’
- Copula removal e.g. ‘です’ → ‘’
- Pronoun substitution e.g. ‘私’ → ‘俺’

A.4 Full Mono-Pair Results

Table 5 shows our full results from the mono-pair model experiments. In addition to CONSIST and BLEU, we also obtain a MATCH score, based on average percentage the output sentence matches the reference and NUM, the average number of different translations generated per set of semantically-equivalent NAV perturbations.

A.5 Multi-pair Model Results

Results from multi-pair experiments in Section 4.2 in Figure 9.

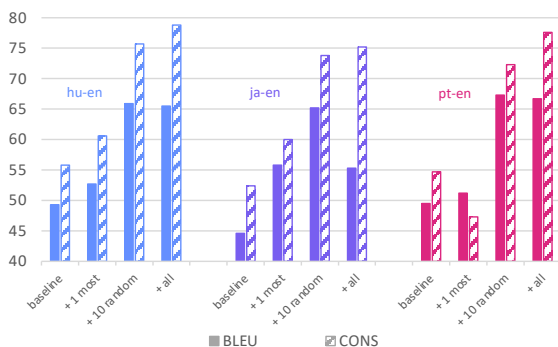


Figure 9: As in the mono-pair translation experiments, we see how ‘+ 1 most’ generally shows improvements over a baseline and ‘+ 10 random’ shows how exposing the model to NAV perturbations further increases both BLEU and CONSIST scores. While including all NAV perturbations continues to improve CONSIST, we see a trade-off in BLEU.

A.6 Quantitative Error

Getting clear quantitative descriptions for the frequency of specific types of errors seen in baseline models is difficult, but we provide some more information about the analysis. We find the types of errors mentioned (issues with dropped pronouns, uncommon synonyms, gender ambiguity, etc.) fairly frequently. Skimming through several prompts, we see these types of erroneous output many times.

For example, in a prompt with 92 NAV perturbations in Japanese, 52 of the outputs from the baseline model were incorrect. With a reference “It is a kind of tomato”, erroneous outputs include “Type of tomatoes” “It is like tomatoes” “Something about tomatoes” and “One sort of tomato”; correct outputs include “It is a variety of tomato” “This is a type of tomato” and “It is a tomato species”. Some of these errors are language specific, since NAV is very dependent on the two languages involved, but as several languages share

hu-en cond	#	Quality		Consistency	
		BLEU	MAT.	CONS.	NUM
base	0	28.2	11.1	43.1	21.3
	1	37.4	22.3	44.8	20.5
	2	42.3	26.7	51.3	16.5
most	10	45.4	29.6	58.0	12.4
	1	46.3	29.8	56.3	12.9
	2	46.9	32.0	58.3	11.9
rand	10	46.6	32.3	62.9	10.0
	1	47.2	30.6	56.4	13.0
	2	48.1	31.4	58.8	11.3
least	10	46.9	31.2	62.8	9.9
	all	43.3	26.4	64.4	8.4

ja-en cond	#	Quality		Consistency	
		BLEU	MAT.	CONS.	NUM
base	0	29.0	14.5	39.5	46.8
	1	45.7	27.3	48.9	29.4
	2	45.6	28.5	49.6	28.1
most	10	50.5	33.6	60.5	16.2
	1	47.9	30.8	54.3	21.6
	2	51.0	33.2	57.3	18.1
rand	10	50.0	33.0	60.3	15.2
	1	49.4	31.7	55.1	21.2
	2	48.8	32.2	56.8	18.6
least	10	50.1	33.1	61.7	13.9
	all	40.8	22.6	66.4	9.1

pt-en cond	#	Quality		Consistency	
		BLEU	MAT.	CONS.	NUM
base	0	38.4	12.0	35.3	34.3
	1	42.0	14.1	34.2	34.0
	2	43.4	16.2	37.5	29.8
most	10	52.8	26.1	53.3	16.9
	1	52.4	26.3	51.6	18.5
	2	54.8	26.3	55.8	15.7
rand	10	56.9	29.0	59.2	13.0
	1	53.8	27.9	52.8	18.2
	2	55.7	29.4	55.3	15.3
least	10	56.5	30.0	59.3	12.8
	all	53.0	26.9	63.0	10.0

Table 5: Results for translation quality and consistency for three language directions under different fine-tuning conditions. Including additional perturbations increases test consistency consistently.

properties, relevant errors can be seen across languages (e.g. pronoun drop in Japanese, Korean and Portuguese)

B Additional Evaluations

B.1 Out-of-Domain Performance of NAV-robust Models

To further investigate behavior of NAV-robust models, we run experiments to see the effect on performance in OOD tasks. For {hu,ja,pt}-en we use a held-out test split of the Tatoeba corpus. Our results are shown in Figure 10. For {ko,vi}-en we use the official OPUS-100 test set with results shown in Figure 11.

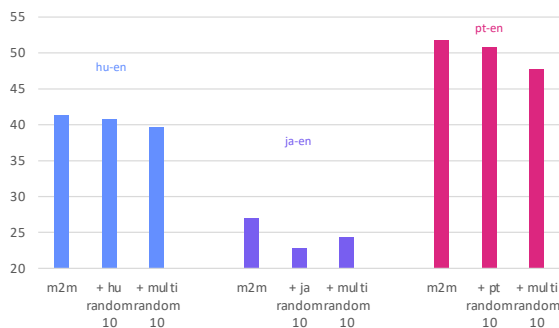


Figure 10: BLEU scores on held-out Tatoeba test sets. Baseline models perform better than NAV-robust models.

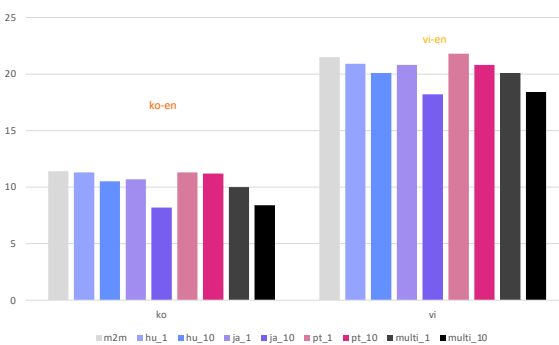


Figure 11: BLEU scores on held-out OPUS-100 test sets. Baseline models perform better than most NAV-robust models.

From these experiments, we see that generally NAV-robustification slightly worsens OOD performance. While these results appear negative, they are not too surprising considering common trends in fine-tuning and robustification (He et al., 2021). Often, fine-tuning can cause some forgetting and models with higher robustness perform worse when evaluated on original, un-noisy input. Correcting

this widespread behavior is beyond the scope of this work.

While it’s clear that improving NAV robustness has shown a decrease in ‘regular’ MT BLEU, there are many differences between our NAV evaluation settings and our experiment settings in this section. For example, these experiments are not just 1) out of domain, they also 2) no longer have a robustness component to them because we aren’t evaluating the models on several perturbations (NAV or otherwise) of input sentences. The decrease in BLEU may be due to either or both of these properties, which our next subsections attempt to disentangle.

B.2 NAV-robust Models on Other Types of Noise

Our next set of experiments in this section looks at measuring NAV-robustness transfer. By this we seek to answer the question *does a model fine-tuned deliberately for NAV robustness exhibit robustness to non-NAV-related noise?* For this, we use an existing MT test set designed to include noisy input text to challenge models’ robustness, MTNT (Machine Translation of Noisy Text) (Michel and Neubig, 2018).

We simply evaluate our baseline and select NAV-tuned models on the ja-en test split of MTNT. Our results are shown in Figure 12. We see that our baseline performs better than our NAV-tuned models, suggesting zero-shot transfer of robustness may not be possible in this way. The NAV perturbations which our models were trained to be robust to do not overlap with many of the types of noise in MTNT, which resembles less-standard ‘internet-speech’.

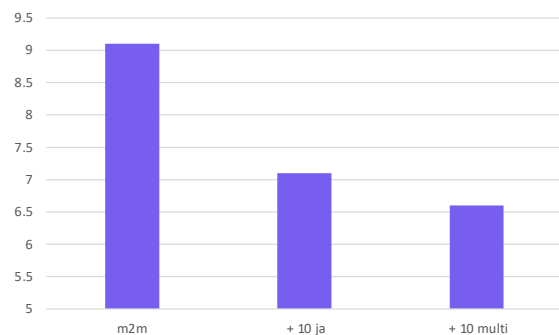


Figure 12: NAV-robust models perform worse than a baseline model in MTNT. MTNT is a difficult task (highlighted by the very low BLEU scores) and includes greater and more varied noise than within the NAV boundaries.

Again, we have another problem with this evaluation in that the models are not only tested against new classes of perturbations but also in a new domain. One way of testing performance on new classes of perturbation but staying in-domain is to use our synthetic augmentation scripts to create a synthetic robustness test set.

We use the same scripts but apply them to the test splits of {hu,ja,pt} STAPLE. In this way, we can probe transferability of NAV robustness by evaluating our NAV-robust models on these synthetic test sets. We also evaluate the models fine-tuned on synthetic data for comparison. Results are shown in Figure 13.

In general, the synthetic-tuned models perform better on these synthetic test evaluations. However, the organic NAV-tuned models do exhibit some improved robustness compared to a baseline model. The NAV perturbations these models are fine-tuned on rarely overlap with the types of synthetic perturbations performed on the new {hu,pt}-en sets, explaining the only slight transferability. Organic NAV perturbations help more on our synthetic ja-en test set, likely because the ja-en test set attempts to synthetically create NAV perturbations, thus imitating the organic fine-tuning data more closely.

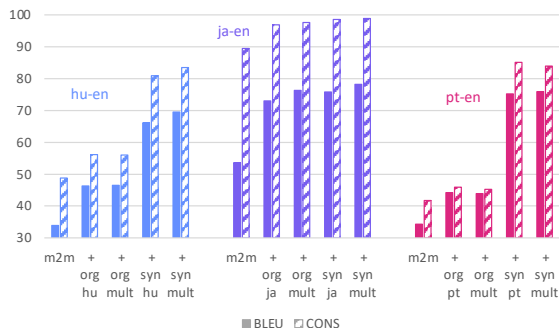


Figure 13: BLEU and CONSIST scores for organic vs. synthetic perturbations during fine-tuning on synthetic STAPLE test sets. Synthetic fine-tuning unsurprisingly improves synthetic test performance, but we also see how organic NAV fine-tuning performs fairly well on rule-based synthetic perturbation which aims to emulate NAV (ja-en) compared to random character-alteration synthetic test sets ({hu,pt}-en).

B.3 In-Domain MT

One possible source of confusion in our NAV robustness evaluation could be the fact that it is abnormal to have hundreds of test examples with the same reference translation. Perhaps BLEU is not as reliable in such conditions. To account for this, we

calculate BLEU on in-domain data by using our ‘1-most’ and ‘1-least’ STAPLE test sets. This returns the task to single-pair examples while focusing on common and uncommon examples, respectively.

The results for ‘1-most’ are shown in Figures 14 and 15. In our NAV experiments from Section 4, the “10 random” models outperform the “1 most” models, but here we see the opposite. This is not surprising as the ‘1-most’ fine-tuning is most similar to this ‘1-most’ evaluation. Also, this evaluation no longer contains a robustness aspect as each example pair uses a commonly-occurring source sentence.

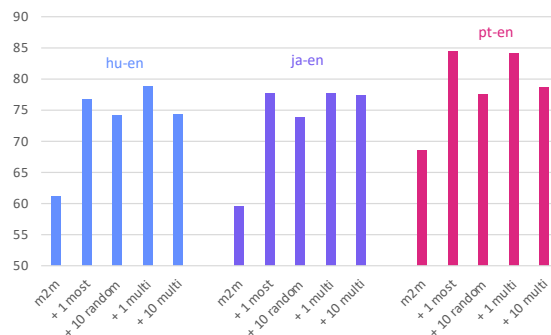


Figure 14: BLEU scores evaluating on 1-most STAPLE test sets. 1-most fine-tuning improves performance the most.

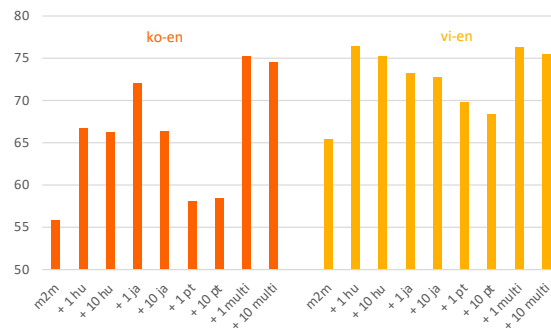


Figure 15: BLEU scores evaluating on 1-most STAPLE test sets for 0-shot language transfer. 1-most fine-tuning improves performance over baseline and 10-random models. We see similar language differences with ja-en improving ko-en more compared to hu-en with vi-en.

This raises the concern that it could be, in fact, that 10-random has higher BLEU on our robustness experiments because of the property of having hundreds of variations for each example. To disentangle this possible explanation, we also evaluate our models on a “1-least” STAPLE set. We source the *least* common source sentence for each example. This helps ensure the test inputs will exhibit

more difficult NAV perturbations, thus serving as a better NAV-robustness test set without the added property of having several variations from each translation pair. The results are shown in Figures 16 and 17.

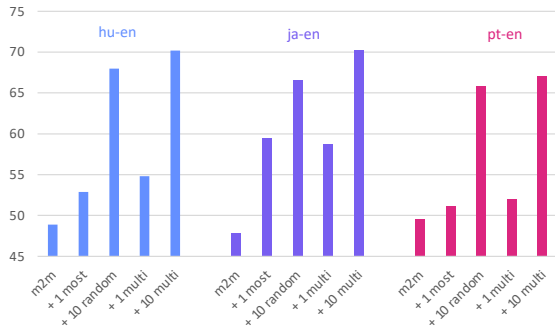


Figure 16: BLEU scores evaluating on 1-least STAPLE test sets. 10-random fine-tuning improves performance the most.

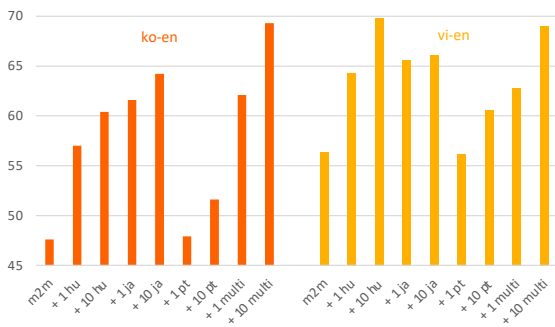


Figure 17: BLEU scores evaluating on 1-least STAPLE test sets for 0-shot language transfer. 10-random fine-tuning improves performance over baseline and 1-most models. We see language differences with ja-en improving ko-en more compared to hu-en with vi-en.

Our results now reflect those from our robustness experiments with “10-random” fine-tuning consistently performing better than “1-most”. This suggests that the models showing highest NAV robustness by our evaluations also show higher robustness in a more standard MT evaluation on NAV-noisy data.