

# Lost in Machine Translation: A Method to Reduce Meaning Loss

Reuben Cohn-Gordon

Stanford

reubencg@stanford.edu

Noah D. Goodman

Stanford

ngoodman@stanford.edu

## Abstract

A desideratum of high-quality translation systems is that they preserve meaning, in the sense that two sentences with different meanings should not translate to one and the same sentence in another language. However, state-of-the-art systems often fail in this regard, particularly in cases where the source and target languages partition the “meaning space” in different ways. For instance, “I cut my finger.” and “I cut my finger off.” describe different states of the world but are translated to French (by both *Fairseq* and *Google Translate*) as “Je me suis coupé le doigt.”, which is ambiguous as to whether the finger is detached. More generally, translation systems are typically many-to-one (non-injective) functions from source to target language, which in many cases results in important distinctions in meaning being lost in translation. Building on Bayesian models of informative utterance production, we present a method to define a less ambiguous translation system in terms of an underlying pre-trained neural sequence-to-sequence model. This method increases injectivity, resulting in greater preservation of meaning as measured by improvement in cycle-consistency, without impeding translation quality (measured by BLEU score).

## 1 Many-to-One Translations

Languages differ in what meaning distinctions they must mark explicitly. As such, translations risk mapping from a form in one language to a more ambiguous form in another. For example, the definite (1) and indefinite (2) both translate (under *Fairseq* and *Google Translate*) to (3) in French, which is ambiguous in definiteness.

*The animals run fast.* (1)

*Animals run fast.* (2)

*Les animaux courent vite* (3)




Figure 1: State-of-the-art neural image captioner  $S_0^{\text{SENT}}$  loses a meaning distinction which *informative* model  $S_1^{\text{SENT-IP}}$  preserves.

**Survey** To evaluate the nature of this problem, we explored a corpus<sup>1</sup> of 500 pairs of distinct English sentences which map to a single German one (the evaluation language in section 2.3). We identify a number of common causes for the many-to-one maps. Two frequent types of verbal distinction lost when translating to German are tense (54 pairs, e.g. “...others {were, have been} introduced.”) and modality (16 pairs, e.g. “...prospects for this year {could, might} be better.”), where German “können” can express both epistemic and ability modality, distinguished in English with “might” and “could” respectively. Owing to English’s large vocabulary, lexical difference in verb (31 pairs, e.g. “arise” vs. “emerge”), noun (56 pairs, e.g. “mystery” vs. “secret”), adjective (47 pairs, e.g. “unaffected” vs. “untouched”) or deictic/pronoun (32 pairs, usually “this” vs “that”) are also common. A large number of the pairs differ instead either orthographically, or in other ways that do not correspond to a clear semantic distinc-

<sup>1</sup>Generated by selecting short sentences from the Brown corpus (Kučera and Francis, 1967), translating them to German, and taking the best two candidate translations back into English, if these two themselves translate to a single German sentence. Translation in both directions was done with Fairseq.

$A$	He is wearing glasses.
$B$	He wears glasses.
$S_0^{\text{SNT}}(A)$	Er trägt eine Brille.
$S_0^{\text{SNT}}(B)$	Er trägt eine Brille .
$S_1^{\text{SNT-IP}}(A)$	Er trägt jetzt eine Brille.
$S_1^{\text{SNT-IP}}(B)$	Er hat eine Brille.

Figure 2: Similar to Figure 1,  $S_0^{\text{SNT}}$  collapses two English sentences into a single German one, whereas  $S_1^{\text{SNT-IP}}$  distinguishes the two in German.

tion (e.g. “She had  $\{taken, made\}$  a decision.”).

**Our approach** While languages differ in what distinctions they are *required* to express, all are usually capable of expressing any given distinction when desired. As such, meaning loss of the kind discussed above is, in theory, avoidable. To this end, we propose a method to reduce meaning loss by applying the Rational Speech Acts (RSA) model of an informative speaker to translation. RSA has been used to model natural language *pragmatics* (Goodman and Frank, 2016), and recent work has shown its applicability to image captioning (Andreas and Klein, 2016; Vedantam et al., 2017; Mao et al., 2016), another sequence-generation NLP task. Here we use RSA to define a translator which reduces many-to-one mappings and consequently meaning loss, in terms of a pretrained neural translation model. We introduce a strategy for performing inference efficiently with this model in the setting of translation, and show gains in *cycle-consistency*<sup>2</sup> as a result. Moreover, we obtain improvements in translation quality (BLEU score), demonstrating that the goal of meaning preservation directly yields improved translations.

## 2 Meaning Preservation as Informativity

In the RSA framework, speakers and listeners, modeled as Bayesian agents, reason about each other in a nested fashion. We refer to listeners and speakers which do not reason about another agent as  $L_0$  and  $S_0$  respectively, and an agent which reasons about another agent as  $L_1$  or  $S_1$ . For instance, an informative speaker model  $S_1$  is given a state

<sup>2</sup>Formally, say that a pair of functions  $f : A \rightarrow B$ ,  $g : B \rightarrow A$  is cycle-consistent if  $g \cdot f = id$ , the identity function. If  $f$  is *not* one-to-one, then  $(f, g)$  is not cycle-consistent. (Note however that when  $A$  and  $B$  are infinite, the converse does not hold: even if  $f$  and  $g$  are both one-to-one,  $(f, g)$  need not be cycle-consistent, since many-to-one maps between infinite sets are not necessarily bijective.)

$w \in W$ , and chooses an utterance  $u \in U$  to convey  $w$  to  $S_1$ ’s model of a listener. By contrast,  $S_0$  chooses utterances without a listener model in mind — its behavior might be determined by a semantics, or in our case, by a pretrained neural model.

For translation, the state space  $W$  is a set of source language sentences (sequences of words in the language), while  $U$  is a set of target language sentences.  $S_1$ ’s goal is to choose a translation  $u$  which allows a listener to pick out the source sentence  $w$  from among the set of distractors. This informative behavior discourages many-to-one maps that a non-informative translation model  $S_0$  might allow.

**$S_0$  Model** BiLSTMs with attention (Bahdanau et al., 2014), and more recently CNNs (Gehring et al., 2016) and entirely attention based models (Vaswani et al., 2017) constitute the state-of-the-art architectures in neural machine translation. All of these systems, once trained end-to-end on aligned data, can be viewed as a conditional distribution<sup>3</sup>  $S_0^{\text{WD}}(wd|w, c)$ , for a word  $wd$  in the target language, a source language sentence  $w$ , and a partial sentence  $c$  in the target language.  $S_0^{\text{WD}}$  yields a distribution  $S_0^{\text{SNT}}$  over full sentences<sup>4</sup>:

$$S_0^{\text{SNT}}(u|w, c) = \prod_t S_0^{\text{WD}}(u[t]|w, c + u[:t]) \quad (4)$$

$S_0^{\text{SNT}}$  returns a distribution over continuations of  $c$  into full target language sentences<sup>5</sup>. To obtain a sentence from  $S_0^{\text{SNT}}$  given a source language sentence  $s$ , one can greedily choose the highest probability word from  $S_0^{\text{WD}}$  at each timestep, or explore a beam of possible candidates. We implement  $S_0^{\text{WD}}$  (in terms of which all our other models are defined) using *Fairseq*’s publicly available<sup>6</sup> pretrained Transformer models for English-German, and for German-English train a CNN using *Fairseq*.

### 2.1 Explicit Distractors

We first describe a sentence level, *globally pragmatic* model  $S_1^{\text{SNT-GP}}$  for the simple case where

<sup>3</sup>We use  $S_{0/1}^{\text{WD}}$  and  $S_{0/1}^{\text{SNT}}$  respectively to distinguish word and sentence level speaker models

<sup>4</sup>Python list indexing conventions are used, “+” means concatenation of list to element or list

<sup>5</sup>In what follows, we omit  $c$  when it is empty, so that  $S_0^{\text{SNT}}(u|w)$  is the probability of sentence  $u$  given  $w$

<sup>6</sup><https://github.com/pytorch/fairseq>

a source language sentence needs to be distinguished from a presupplied distractor<sup>7</sup> (as in the pairs shown in figures (2) and (1)). We use this model as a stepping stone to one which requires an input sentence in the source language only, and no distractors. We begin by defining a listener  $L_1^{\text{SNT}}$ , which receives a target language sentence  $u$  and infers which sentence  $w \in W$  (a presupplied set such as the pair (1) and (2)) would have resulted in the pretrained neural model  $S_0^{\text{SNT}}$  producing  $u$ :

$$L_1^{\text{SNT}}(w|u) \propto \frac{S_0^{\text{SNT}}(u|w)}{\sum_{w' \in W} S_0^{\text{SNT}}(u|w')} \quad (5)$$

This allows  $S_1^{\text{SNT-GP}}$  to be defined in terms of  $L_1^{\text{SNT}}$ , where  $U$  is the set of all possible target language sentences<sup>8</sup>:

$$S_1^{\text{SNT-GP}}(u|w) = \frac{S_0^{\text{SNT}}(u|w)L_1^{\text{SNT}}(w|u)^\alpha}{\sum_{u' \in U} S_0^{\text{SNT}}(u'|w)L_1^{\text{SNT}}(w|u')^\alpha} \quad (6)$$

The key property of this model is that, for  $W = \{A, B\}$ , when translating  $A$ ,  $S_1^{\text{SNT-GP}}$  prefers translations of  $A$  that are unlikely to be good translations of  $B$ . So for pairs like (1) and (2),  $S_1^{\text{SNT-GP}}$  is compelled to produce a translation for the former that reflects its difference from the latter, and vice versa.

**Inference** Since  $U$  is an infinite set, exactly computing the most probable utterance under  $S_1^{\text{SNT-GP}}(\cdot|w)$  is intractable. Andreas and Klein (2016) and Mao et al. (2016) perform approximate inference by sampling the subset of  $U$  produced by a beam search from  $S_0^{\text{SNT}}$ . Vedantam et al. (2017) and Cohn-Gordon et al. (2018) employ a different method, using an *incremental* model  $S_1^{\text{SNT-IP}}$  as an approximation of  $S_1^{\text{SNT-GP}}$  on which inference can be tractably performed.

$S_1^{\text{SNT-IP}}$  considers informativity not over the whole set of utterances, but instead at each decision of the next word in the target language sentence. For this reason, the incremental method avoids the problem of lack of beam diversity encountered when subsampling from  $S_0^{\text{SNT}}$ , which

<sup>7</sup>Implementations for all models are available to <https://github.com/reubenharry/pragmatic-translation>

<sup>8</sup> $\alpha$  is a hyperparameter of  $S_1^{\text{SNT-GP}}$ ; as it increases, the model cares more about being informative and less about producing a reasonable translation.

becomes especially bad when producing long sequences, as is often the case in translation.  $S_1^{\text{SNT-IP}}$  is defined as the product of informative decisions, specified by  $S_1^{\text{WD}}$  (itself defined in terms of  $L_1^{\text{WD}}$ ), which are defined analogously to (6) and (5).

$$L_1^{\text{WD}}(w|wd, c) \propto S_0^{\text{WD}}(wd|w, c) \quad (7)$$

$$S_1^{\text{WD}}(wd|w, c) \propto S_0^{\text{WD}}(wd|w, c) * L_1^{\text{WD}}(w|wd, c)^\alpha \quad (8)$$

$$S_1^{\text{SNT-IP}}(u|w, c) = \prod_t S_1^{\text{WD}}(u[t]|w, c + u[:t]) \quad (9)$$

**Examples**  $S_1^{\text{SNT-IP}}$  is able to avoid many-to-one mappings by choosing more informative translations. For instance, its translation of (1) is “Ces animaux courent vite” (*These animals run fast.*). See figures (1) and (2) for other examples of many-to-one mappings under  $S_0^{\text{SNT}}$  avoided by  $S_1^{\text{SNT-IP}}$ .

## 2.2 Avoiding Explicit Distractors

While  $S_1^{\text{SNT-IP}}$  can disambiguate between pairs of sentences, it has two shortcomings. First, it requires one (or more) distractors to be provided, so translation is no longer fully automatic. Second, because the distractor set  $W$  consists of only a pair (or finite set) of sentences,  $S_1^{\text{SNT-IP}}$  only cares about being informative up to the goal of distinguishing between these sentences. Intuitively, total meaning preservation is achieved by a translation which distinguishes the source sentence  $w$  from every other sentence in the source language which differs in meaning.

Both of these problems can be addressed by introducing a new “cyclic” globally pragmatic model  $S_1^{\text{SNT-CGP}}$  which reasons not about  $L_1^{\text{SNT}}$  but about a pretrained translation model from target language to source language, which we term  $L_0^{\text{SNT}}$ .

$$S_1^{\text{SNT-CGP}}(u|w) \propto S_0^{\text{SNT}}(u|w)L_0^{\text{SNT}}(w|u)^\alpha \quad (10)$$

$S_1^{\text{SNT-CGP}}$  is like  $S_1^{\text{SNT-GP}}$ , but its goal is to produce a translation which allows a listener model (now  $L_0^{\text{SNT}}$ ) to infer the original sentence, not among a small set of presupplied possibilities, but among *all source language sentences*. As such, an optimal translation  $u$  of  $w$  under  $S_1^{\text{SNT-CGP}}$  has high probability of being generated by  $S_0^{\text{SNT}}$  and

high probability of being back-translated to  $w$  by  $L_0^{\text{SNT}}$ .  $S_1^{\text{SNT-CGP}}$  is very closely related to reconstruction methods, e.g. (Tu et al., 2017).

**Incremental Model** Exact inference is again intractable, though as with  $S_1^{\text{SNT-GP}}$ , it is possible to approximate by subsampling from  $S_0^{\text{SNT}}$ . This is very close to the approach taken by (Li et al., 2016), who find that reranking a set of outputs by probability of recovering input “dramatically decreases the rate of dull and generic responses.” in a question-answering task. However, because the subsample is small relative to  $U$ , they use this method in conjunction with a diversity increasing decoding algorithm.

As in the case with explicit distractors, we instead opt for an incremental model, now  $S_1^{\text{SNT-CIP}}$  which approximates  $S_1^{\text{SNT-CGP}}$ . The definition of  $S_1^{\text{SNT-CIP}}$  (12) is more complex than the incremental model with explicit distractors ( $S_1^{\text{SNT-IP}}$ ) since  $L_0^{\text{WD}}$  must receive complete sentences, rather than partial ones like  $L_1^{\text{WD}}$ . As such, we need to marginalize over continuations  $k$  of partial sentences in the target language:

$$S_1^{\text{WD-C}}(wd|w, c) \propto S_0^{\text{WD}}(wd|w, c) * \sum_k (L_0^{\text{SNT}}(w|c + wd + k) S_0^{\text{SNT}}(k|w, c + wd)) \quad (11)$$

$$S_1^{\text{SNT-CIP}}(u|w, c) = \prod_t S_1^{\text{WD-C}}(u[t]|w, c + u[:t]) \quad (12)$$

Since the sum over continuations of  $c$  in (11) is intractable to compute exactly, we approximate it by a single continuation, obtained by greedily unrolling  $S_0^{\text{SNT}}$ . The whole process of generating a new word  $wd$  of the translation from a sequence  $c$  and a source language sentence  $w$  is as follows: first use  $S_0^{\text{WD}}$  to generate a set of candidates for the next word (in practice, we only consider two, for efficiency). For each of these, use  $S_0^{\text{SNT}}$  to greedily unroll a full target language sentence from  $c + wd$ , namely  $c + wd + k$ , and rank each  $wd$  by the probability  $L_0^{\text{SNT}}(w|c + wd + k)$ .

### 2.3 Evaluating the Informative Translator

Our objective is to improve meaning preservation without detracting from translation quality in other regards (e.g. grammaticality). We conduct our evaluations on English to German translation, making use of publicly available

pre-trained English-German and German-English Fairseq models. The pragmatic model we evaluate is  $S_1^{\text{SNT-CIP}}$  since, unlike  $S_1^{\text{SNT-IP}}$ , it is not necessary to hand-supply a distractor set of source language sentences.

An example of the behavior of  $S_1^{\text{SNT-CIP}}$  and  $S_0^{\text{SNT}}$  on of our test sentences is shown below;  $S_0^{\text{SNT}}$  is able to preserve the phrase “world’s eyes”, which  $S_0^{\text{SNT}}$  translates merely as “world”:

- Source sentence: Isolation keeps the world’s eyes off Papua.
- Reference translation: Isolation hält die Augen der Welt fern von Papua.
- $S_0^{\text{SNT}}$ : Die Isolation hält die Welt von Papua fern.
- $S_1^{\text{SNT-CIP}}$ : Die Isolation hält die Augen der Welt von Papua fern.

We use cycle-consistency as a measure of meaning preservation, since the ability to recover the original sentence requires meaning distinctions not to be collapsed. In evaluating cycle-consistency, it is important to use a separate target-source translation mechanism than the one used to define  $S_1^{\text{SNT-CIP}}$ . Otherwise, the system has access to the model which evaluates it and may improve cycle-consistency without producing meaningful target language sentences. For this reason, we translate German sentences (produced by  $S_0^{\text{SNT}}$  or  $S_1^{\text{SNT-CIP}}$ ) back to English with *Google Translate*. To measure cycle-consistency, we use the BLEU metric (implemented with sacreBLEU (Post, 2018)), with the original sentence as the reference.

However, this improvement of cycle consistency, especially with a high value of  $\alpha$ , may come at the cost of translation quality. Moreover, it is unclear whether BLEU serves as a good metric for evaluating sentences of a single language. To further ensure that translation quality is not compromised by  $S_1^{\text{SNT-CIP}}$ , we evaluate BLEU scores of the German sentences it produces. This requires evaluation on a corpus of aligned sentences, unlike the sentences collected from the Brown corpus in section 1<sup>9</sup>.

<sup>9</sup>While we find that  $S_1^{\text{SNT-CIP}}$  improves cycle-consistency for the Brown corpus over  $S_0^{\text{SNT}}$ , we have no way to establish whether this comes at the cost of translation quality.



Model	Cycle	Translate
$S_0^{\text{SNT}}$	43.35	37.42
$S_1^{\text{SNT-CIP}}$	<b>47.34</b>	<b>38.29</b>

Table 1: BLEU score on cycle-consistency and translation for WMT, across baseline and informative models. Greedy unrolling and  $\alpha = 0.1$

We perform both evaluations (cycle-consistency and translation) on 750 sentences<sup>10</sup> of the 2018 English-German WMT News test-set.<sup>11</sup> We use greedy unrolling in all models (using beam search is a goal for future work). For  $\alpha$  (which represents the trade-off between informativity and translation quality) we use 0.1, obtained by tuning on validation data.

**Results** As shown in table (1),  $S_1^{\text{SNT-CIP}}$  improves over  $S_0^{\text{SNT}}$  not only in cycle-consistency, but in translation quality as well. This suggests that the goal of preserving information, in the sense defined by  $S_1^{\text{SNT-CGP}}$  and approximated by  $S_1^{\text{SNT-CIP}}$ , is important for translation quality.

### 3 Conclusions

We identify a shortcoming of state-of-the-art translation systems and show that a version of the RSA framework’s informative speaker  $S_1$ , adapted to the domain of translation, alleviates this problem in a way which improves not only cycle-consistency but translation quality as well. The success of  $S_1^{\text{SNT-CIP}}$  on two fairly similar languages raises the question of whether improvements will increase for more distant language pairs, in which larger scale differences exist in what information is obligatorily represented - this is a direction for future work.

### Acknowledgements

Thanks to the reviewers for their substantive comments, and to Daniel Fried and Jacob Andreas for many helpful discussions during the development of this project.

<sup>10</sup>Our implementation of  $S_1^{\text{SNT-CIP}}$  was not efficient, and we could not evaluate on more sentences for reasons of time.

<sup>11</sup><http://www.statmt.org/wmt18/translation-task.html>

### References

- Jacob Andreas and Dan Klein. 2016. Reasoning about pragmatics with neural listeners and speakers. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182. Association for Computational Linguistics.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Reuben Cohn-Gordon, Noah Goodman, and Christopher Potts. 2018. Pragmatically informative image captioning with character-level inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 439–443. Association for Computational Linguistics.
- Jonas Gehring, Michael Auli, David Grangier, and Yann N Dauphin. 2016. A convolutional encoder model for neural machine translation. *arXiv preprint arXiv:1611.02344*.
- Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11):818–829.
- Henry Kučera and Winthrop Nelson Francis. 1967. *Computational analysis of present-day American English*. Dartmouth Publishing Group.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. A simple, fast diverse decoding algorithm for neural generation. *arXiv preprint arXiv:1611.08562*.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20.
- Matt Post. 2018. A call for clarity in reporting bleu scores. *arXiv preprint arXiv:1804.08771*.
- Zhaopeng Tu, Yang Liu, Lifeng Shang, Xiaohua Liu, and Hang Li. 2017. Neural machine translation with reconstruction. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- 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*, pages 5998–6008.
- Ramakrishna Vedantam, Samy Bengio, Kevin Murphy, Devi Parikh, and Gal Chechik. 2017. Context-aware captions from context-agnostic supervision. In *Computer Vision and Pattern Recognition (CVPR)*, volume 3.