

# Automatic Reference-Based Evaluation of Pronoun Translation Misses the Point

Liane Guillou\*

University of Edinburgh  
School of Informatics  
Scotland, United Kingdom  
lguillou@inf.ed.ac.uk

Christian Hardmeier\*

Uppsala University  
Dept. of Linguistics & Philology  
Uppsala, Sweden  
christian.hardmeier@lingfil.uu.se

## Abstract

We compare the performance of the APT and AutoPRF metrics for pronoun translation against a manually annotated dataset comprising human judgements as to the correctness of translations of the PROTEST test suite. Although there is some correlation with the human judgements, a range of issues limit the performance of the automated metrics. Instead, we recommend the use of semi-automatic metrics and test suites in place of fully automatic metrics.

## 1 Introduction

As the general quality of machine translation (MT) increases, there is a growing interest in improving the translation of specific linguistic phenomena. A case in point that has been studied in the context of both statistical (Hardmeier, 2014; Guillou, 2016; Loáiciga, 2017) and neural MT (Bawden et al., 2017; Voita et al., 2018) is that of pronominal anaphora. In the simplest case, translating anaphoric pronouns requires the generation of corresponding word forms respecting the grammatical constraints on agreement in the target language, as in the following English-French example, where the correct form of the pronoun in the second sentence varies depending on which of the (equally correct) translations of the word *bicycle* was used in the first:

- (1) a. I have a bicycle. **It** is red.
- b. J’ai un vélo. **Il** est rouge. [ref]
- c. J’ai une bicyclette. **Elle** est rouge. [MT]

However, the problem is more complex in practice because there is often no 1 : 1 correspondence between pronouns in two languages. This is easily demonstrated at the corpus level by observing that the number of pronouns varies significantly across languages in parallel texts (Mitkov

and Barbu, 2003), but it tends to be difficult to predict in individual cases.

In general MT research, significant progress was enabled by the invention of automatic evaluation metrics based on reference translations, such as BLEU (Papineni et al., 2002). Attempting to create a similar framework for efficient research, researchers have proposed automatic reference-based evaluation metrics specifically targeting pronoun translation: AutoPRF (Hardmeier and Federico, 2010) and APT (Miculicich Werlen and Popescu-Belis, 2017). We study the performance of these metrics on a dataset of English-French translations and investigate to what extent automatic evaluation based on reference translations provides insights into how well an MT system handles pronouns. Our analysis clarifies the conceptual differences between AutoPRF and APT, uncovering weaknesses in both metrics, and investigates the effects of the alignment correction heuristics used in APT. By using the fine-grained PROTEST categories of pronoun function, we find that the accuracy of the automatic metrics varies across pronouns of different functions, suggesting that certain linguistic patterns are captured better in the automatic evaluation than others. We argue that fully automatic wide-coverage evaluation of this phenomenon is unlikely to drive research forward, as it misses essential parts of the problem despite achieving some correlation with human judgements. Instead, semi-automatic evaluation involving automatic identification of correct translations with high precision and low recall appears to be a more achievable goal. Another more realistic option is a test suite evaluation with a very limited scope.

## 2 Pronoun Evaluation Metrics for MT

Two reference-based automatic metrics of pronoun translation have been proposed in the literature.

\*Both authors contributed equally.

The first (Hardmeier and Federico, 2010) is a variant of precision, recall and F-score that measures the overlap of pronouns in the MT output with a reference translation. It lacks an official name, so we refer to it as *AutoPRF* following the terminology of the DiscoMT 2015 shared task (Hardmeier et al., 2015). The scoring process relies on a word alignment between the source and the MT output, and between the source and the reference translation. For each input pronoun, it computes a *clipped count* (Papineni et al., 2002) of the overlap between the aligned tokens in the reference and the MT output. The clipped count of a given word is defined as the number of times it occurs in the MT output, limited by the number of times it occurs in the reference translation. The final metric is then calculated as the precision, recall and F-score based on these clipped counts.

Miculicich Werlen and Popescu-Belis (2017) propose a metric called *Accuracy of Pronoun Translation (APT)* that introduces several innovations over the previous work. It is a variant of accuracy, so it counts, for each source pronoun, whether its translation can be considered correct, without considering multiple alignments. Since word alignment is problematic for pronouns, the authors propose an heuristic procedure to improve alignment quality. Finally, it introduces the notion of pronoun equivalence, assigning partial credit to pronoun translations that differ from the reference translation in specific ways deemed to be acceptable. In particular, it considers six possible cases when comparing the translation of a pronoun in MT output and the reference. The pronouns may be: (1) *identical*, (2) *equivalent*, (3) *different/incompatible*, or there may be no translation in: (4) the MT output, (5) the reference, (6) either the MT output or the reference. Each of these cases may be assigned a weight between 0 and 1 to determine the level of *correctness*.

### 3 The PROTEST Dataset

We study the behaviour of the two automatic metrics using the PROTEST test suite (Guillou and Hardmeier, 2016). The test suite comprises 250 hand-selected personal pronoun tokens taken from the *DiscoMT2015.test* dataset of TED talk transcriptions and translations (Hardmeier et al., 2016) and annotated according to the ParCor guidelines (Guillou et al., 2014). It is structured according to a linguistic typology motivated by work on func-

tional grammar by Dik (1978) and Halliday (2004). Pronouns are first categorised according to their *function*:

- anaphoric*: I have a bicycle. **It** is red.
- event*: He lost his job. **It** was a shock.
- pleonastic*: **It** is raining.
- addressee reference*: **You**'re welcome.

They are then subcategorised according to morphosyntactic criteria, whether the antecedent is a group noun, whether the antecedent is in the same or a different sentence, and whether an addressee reference pronoun refers to one or more specific people (deictic) or to people in general (generic).

Our dataset contains human judgements on the performance of nine MT systems on the translation of the 250 pronouns in the PROTEST test suite. The systems include five submissions to the DiscoMT 2015 shared task on pronoun translation (Hardmeier et al., 2015) – four phrase-based SMT systems AUTO-POSTEDIT (Guillou, 2015), UU-HARDMEIER (Hardmeier et al., 2015), IDIAP (Luong et al., 2015), UU-TIEDEMANN (Tiedemann, 2015), a rule-based system ITS2 (Loáiciga and Wehrli, 2015), and the shared task baseline (also phrase-based SMT). Three NMT systems are included for comparison: LIMSI (Bawden et al., 2017), NYU (Jean et al., 2014), and YANDEX (Voita et al., 2018).

Manual evaluation was conducted using the PROTEST graphical user interface and accompanying guidelines (Hardmeier and Guillou, 2016). The annotators were asked to make judgements (correct/incorrect) on the translations of the pronouns and antecedent heads whilst ignoring the correctness of other words (except in cases where it impacted the annotator's ability to make a judgement). The annotations were carried out by two bilingual English-French speakers, both of whom are native speakers of French. Our human judgements differ in important ways from the human evaluation conducted for the same set of systems at DiscoMT 2015 (Hardmeier et al., 2015), which was carried out by non-native speakers over an unbalanced data sample using a gap-filling methodology. In the gap-filling task annotators are asked to select, from a predefined list (including an uninformative catch-all group "other"), those pronouns that could fill the pronoun translation slot. Unlike in the PROTEST evaluation, the pronoun translations were obscured in the MT output. This avoided priming the annotators with the output of

the candidate translation, but it occasionally caused valid translations to be rejected because they were missed by the annotator.

#### 4 Accuracy versus Precision/Recall

There are three ways in which APT differs from AutoPRF: the scoring statistic, the alignment heuristic in APT, and the definition of pronoun equivalence.

APT is a measure of accuracy: It reflects the proportion of source pronouns for which an acceptable translation was produced in the target. AutoPRF, by contrast, is a precision/recall metric on the basis of clipped counts. Hardmeier and Federico (2010) motivate the use of precision and recall by pointing out that word alignments are not 1 : 1, so each pronoun can be linked to multiple elements in the target language, both in the reference translation and in the MT output. Their metric is designed to account for all linked words in such cases.

To test the validity of this argument, we examined the subset of examples of 8 systems in our English-French dataset<sup>1</sup> giving rise to a clipped count greater than one<sup>2</sup> and found that these examples follow very specific patterns. All 143 cases included exactly one personal pronoun. In 99 cases, the additional matched word was the complementiser *que* ‘that’. In 31 and 4 cases, respectively, it was a form of the auxiliary verbs *avoir* ‘to have’ and *être* ‘to be’. One example matched both *que* and a form of *être*. Two had reflexive pronouns, and one an imperative verb form. With the possible exception of the two reflexive pronouns, none of this seems to be relevant to pronoun correctness. We conclude that it is more reasonable to restrict the counts to a single pronominal item per example. With this additional restriction, however, the recall score of AutoPRF becomes equivalent to a version of APT without equivalent pronouns and alignment correction. We therefore limit the remainder of our study to APT.

#### 5 Effects of Word Alignment

APT includes an heuristic alignment correction procedure to mitigate errors in the word alignment between a source-language text and its translation (reference or MT output). We ran experiments to

<sup>1</sup>Excluding the YANDEX system, which was added later.

<sup>2</sup>A clipped count greater than one for a given pronoun translation indicates that the MT output and the reference translation aligned to this pronoun overlap in more than one token.

Score Alig. corr.	APT-A		APT-B		PRO- TEST
	+	-	+	-	
Reference	1.000	1.000	1.000	1.000	0.920
BASELINE	<b>0.544</b>	<b>0.536</b>	<b>0.574</b>	<b>0.566</b>	0.660
IDIAP	0.496	0.496	0.528	0.528	0.660
UU-TIED.	0.532	0.532	0.562	0.562	0.680
UU-HARD.	0.528	0.520	0.556	0.548	0.636
POSTEDIT	0.492	0.492	0.532	0.532	0.668
ITS2	0.436	0.428	0.462	0.454	0.472
LIMSI	0.364	0.364	0.388	0.388	0.576
NYU	0.424	0.420	0.456	0.452	0.616
YANDEX	<b>0.544</b>	<b>0.536</b>	0.570	0.562	<b>0.796</b>

Table 1: Comparison of APT scores with human judgements over the PROTEST test suite

assess the correlation of APT with human judgements, with and without the alignment correction heuristics.

Table 1 displays the APT results in both conditions and the proportion of pronouns in the PROTEST test suite marked as correctly translated. For better comparison with the PROTEST test suite results, we restricted APT to the pronouns in the test suite. We used two different weight settings:<sup>3</sup> **APT-A** uses weight 1 for *identical* matches and 0 for all other cases. **APT-B** uses weight 1 for *identical* matches, 0.5 for *equivalent* matches and 0 otherwise.

There is little difference in the APT scores when we consider the use of alignment heuristics. This is due to the small number of pronouns for which alignment improvements are applied for most systems (typically 0–12 per system). The exception is the ITS2 system output for which 18 alignment improvements are made. For the following systems we observe a very small increase in APT score for each of the two weight settings we consider, when alignment heuristics are applied: UU-HARDMEIER (+0.8), ITS2 (+0.8), BASELINE (+0.8), YANDEX (+0.8), and NYU (+0.4). However, these small improvements are not sufficient to affect the system rankings. It seems, therefore, that the alignment heuristic has only a small impact on the validity of the score.

To assess differences in correlation with human judgment for pairs of APT settings, we run Williams’s significance test (Williams, 1959; Graham and Baldwin, 2014). The test reveals that differences in correlation between the various configurations of APT and human judgements are not statistically significant ( $p > 0.2$  in all cases).

<sup>3</sup>Personal recommendation by Lesly Miculicich Werlen.

	c1	c2	Pearson	Spearman
With alignment heuristics	1	0	0.848	0.820
Without alignment heuristics	1	0	0.850	0.820
	1	0.5	0.853	0.815
	1	0.5	0.855	0.811

Table 2: Correlation of APT and human judgements

Category	APT Cases			Human Assess.		Disagreement		
	1	2	3	✓	✗	Dis.	Ex.	%
Anaphoric								
intra sbj <i>it</i>	130	13	73	156	60	47/	216	21.8
intra nsbj <i>it</i>	59	1	31	77	14	19/	91	20.9
inter sbj <i>it</i>	104	21	111	142	94	63/	236	26.7
inter nsbj <i>it</i>	21	0	7	8	20	13/	28	<b>46.4</b>
intra <i>they</i>	131	0	95	154	72	37/	226	16.4
inter <i>they</i>	126	0	108	129	105	47/	234	20.1
sg <i>they</i>	57	0	66	83	40	58/	123	<b>47.2</b>
group <i>it/they</i>	47	0	41	64	24	31/	88	35.2
Event <i>it</i>	145	42	94	185	96	60/	281	21.4
Pleonastic <i>it</i>	171	54	52	243	34	46/	277	16.6
Generic <i>you</i>	117	0	70	186	1	69/	187	<b>36.9</b>
Deictic sg <i>you</i>	95	0	47	140	2	45/	142	31.7
Deictic pl <i>you</i>	91	0	7	97	1	6/	98	6.1
Total	1,294	131	802	1,664	563	541/2,227	24.3	

Table 3: Number of pronouns marked as correct/incorrect in the PROTEST human judgements, as identical (1), equivalent (2), and incompatible (3) by APT, and the percentage of disagreements, per category (Disagree [Dis.] / Examples [Ex.])

## 6 Metric Accuracy per Category

Like Miculicich Werlen and Popescu-Belis (2017), we use Pearson’s and Spearman’s correlation coefficients to assess the correlation between APT and our human judgements (Table 2). Although APT does correlate with the human judgements over the PROTEST test suite, the correlation is weaker than that with the DiscoMT gap-filling evaluations reported in Miculicich Werlen and Popescu-Belis (2017). A Williams significance test reveals that the difference in correlation (for those systems common to both studies) is not statistically significant ( $p > 0.3$ ). Table 1 also shows that the rankings induced from the PROTEST and APT scores are rather different. The differences are due to the different ways in which the two metrics define pronoun correctness, and the different sources against which correctness is measured (reference translation vs. human judgement).

We also study how the results of APT (with alignment correction) interact with the categories in PROTEST. We consider a pronoun to be measured as *correct* by APT if it is assigned case 1

(identical) or 2 (equivalent). Likewise, a pronoun is considered *incorrect* if it is assigned case 3 (incompatible). We compare the number of pronouns marked as correct/incorrect by APT and by the human judges, ignoring APT cases in which no judgement can be made: no translation of the pronoun in the MT output, reference or both, and pronouns for which the human judges were unable to make a judgement due to factors such as poor overall MT quality, incorrect word alignments, etc. The results of this comparison are displayed in Table 3.

At first glance, we can see that APT disagrees with the human judgements for almost a quarter (24.3%) of the assessed translations. The distribution of the disagreements over APT cases is very skewed and ranges from 8% for case 1 to 32% for case 2 and 49% for case 3. In other words, APT identifies correct pronoun translations with good precision, but relatively low recall. We can also see that APT rarely marks pronouns as equivalent (case 2).

Performance for anaphoric pronouns is mixed. In general, there are three main problems affecting anaphoric pronouns (Table 4). 1) APT, which does not incorporate knowledge of anaphoric pronoun antecedents, does not consider pronoun-antecedent head agreement so many *valid alternative translations involving personal pronouns* are marked as incompatible (i.e. incorrect, case 3), but as correct by the human judges. Consider the following example, in which the pronoun *they* is deemed correctly translated by the YANDEX system (according to the human judges) as it agrees in number and grammatical gender with the translation of the antecedent *extraits (clips)*. However, the pronoun translation *ils* is marked as incorrect by APT as it does not match the translation in the reference (*elles*).

SOURCE: so what these two **clips** show is not just the devastating consequence of the disease, but **they** also tell us something about the shocking pace of the disease...

YANDEX: donc ce que ces deux **extraits**[masc.,pl.] montrent n’est pas seulement la conséquence dévastatrice de la maladie, mais **ils**[masc. pl.] nous disent aussi quelque chose sur le rythme choquant de la maladie...

REFERENCE: ce que ces deux **vidéos**[fem.,pl.] montrent, ce ne sont pas seulement les conséquences dramatiques de cette maladie, **elles**[fem. pl.] nous montrent aussi la vitesse fulgurante de cette maladie...



2) Substitutions between pronouns are governed by much more complex rules than the simple pronoun equivalence mechanism in APT. For example, the dictionary of pronouns used in APT lists *il* and *ce* as equivalent. However, while *il* can often replace *ce* as a pleonastic pronoun in French, it has a much stronger tendency to be interpreted as anaphoric, rendering pleonastic use unacceptable if there is a salient masculine antecedent in the context. 3) APT does not consider the use of *impersonal pronouns* such as *c'* in place of the feminine personal pronoun *elle* or the plural forms *ils* and *elles*.

Category	V	E	I	O
Anaphoric				
intra-sent. subj. <i>it</i>	22	9	8	8
intra-sent. non-subj. <i>it</i>	16	–	1	2
inter-sent. subj. <i>it</i>	35	6	22	–
inter-sent. non-subj. <i>it</i>	–	–	–	13
intra-sent. <i>they</i>	25	–	3	9
inter-sent. <i>they</i>	22	–	3	22
singular <i>they</i>	40	–	–	18
group <i>it/they</i>	21	–	–	10
<hr/>				
Event <i>it</i>	–	16	–	44
<hr/>				
Pleonastic <i>it</i>	–	11	–	35

V: Valid alternative translation    I: Impersonal translation  
E: Incorrect equivalence        O: Other

Table 4: Common cases of disagreement for anaphoric, pleonastic, and event reference pronouns

As with anaphoric pronouns, APT incorrectly marks some *pleonastic* and *event* translations as equivalent, in disagreement with the human judges. Other common errors arise from 1) the use of alternative translations marked as incompatible (i.e. incorrect) by APT but correct by the human judges, for example *il* (personal) in the MT output when the reference contained the impersonal pronoun *cela* or *ça* (30 cases for pleonastic, 7 for event), or 2) the presence of *il* in both the MT output and reference marked by APT as identical but by the human judges as incorrect (3 cases for pleonastic, 15 event).

Some of these issues could be addressed by incorporating knowledge of pronoun function in the source language, of pronoun antecedents, and of the wider context of the translation surrounding the pronoun. However, whilst we might be able to derive language-specific rules for some scenarios, it would be difficult to come up with more general or language-independent rules. For example, *il* and *ce* can be anaphoric or pleonastic pronouns, but

*il* has a more *referential* character. Therefore in certain constructions that are strongly pleonastic (e.g. clefts) only *ce* is acceptable. This rule would be specific to French, and would not cover other scenarios for the translation of pleonastic *it*. Other issues include the use of pronouns in impersonal constructions such as *il faut* [one must/it takes] in which evaluation of the pronoun requires consideration of the whole expression, or transformations between active and passive voice, where the perspective of the pronouns changes.

## 7 Conclusions

Our analyses reveal that despite some correlation between APT and the human judgements, fully automatic wide-coverage evaluation of pronoun translation misses essential parts of the problem. Comparison with human judgements shows that APT identifies good translations with relatively high precision, but fails to reward important patterns that pronoun-specific systems must strive to generate. Instead of relying on fully automatic evaluation, our recommendation is to emphasise high precision in the automatic metrics and implement semi-automatic evaluation procedures that refer negative cases to a human evaluator, using available tools and methods (Hardmeier and Guillou, 2016). Fully automatic evaluation of a very restricted scope may still be feasible using test suites designed for specific problems (Bawden et al., 2017).

## Acknowledgements

We would like to thank our annotators, Marie Dubremetz and Miryam de Lhoneux, for their many hours of painstaking work, Lesly Miculicich Werlen for providing APT results for the DiscoMT 2015 systems, Elena Voita, Sébastien Jean, Stanislas Lauly and Rachel Bawden for providing the NMT system outputs, and the three anonymous reviewers. The annotation work was funded by the European Association for Machine Translation. The work carried out at The University of Edinburgh was funded by the ERC H2020 Advanced Fellowship GA 742137 SEMANTAX and a grant from The University of Edinburgh and Huawei Technologies. The work carried out at Uppsala University was funded by the Swedish Research Council under grant 2017-930.

## References

- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2017. [Evaluating discourse phenomena in neural machine translation](#). *CoRR*, abs/1711.00513.
- Simon C. Dik. 1978. *Functional Grammar*. Amsterdam, North-Holland.
- Yvette Graham and Timothy Baldwin. 2014. Testing for significance of increased correlation with human judgment. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 172–176, Doha, Qatar. Association for Computational Linguistics.
- Liane Guillou. 2015. Automatic post-editing for the DiscoMT pronoun translation task. In *Proceedings of the Second Workshop on Discourse in Machine Translation*, pages 65–71, Lisbon, Portugal. Association for Computational Linguistics.
- Liane Guillou. 2016. *Incorporating Pronoun Function into Statistical Machine Translation*. Ph.D. thesis, Edinburgh University, Department of Informatics.
- Liane Guillou and Christian Hardmeier. 2016. PROTEST: A test suite for evaluating pronouns in machine translation. In *Proceedings of the Eleventh Language Resources and Evaluation Conference, LREC 2016*, pages 636–643, Portorož, Slovenia.
- Liane Guillou, Christian Hardmeier, Aaron Smith, Jörg Tiedemann, and Bonnie Webber. 2014. ParCor 1.0: A parallel pronoun-coreference corpus to support statistical MT. In *Proceedings of the 9th International Conference on Language Resources and Evaluation, LREC 2014*, pages 3191–3198, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Michael A. K. Halliday. 2004. *An introduction to functional grammar*, 3rd edition. Hodder Arnold London.
- Christian Hardmeier. 2014. *Discourse in Statistical Machine Translation*. Ph.D. thesis, Uppsala University, Department of Linguistics and Philology.
- Christian Hardmeier and Marcello Federico. 2010. Modelling pronominal anaphora in statistical machine translation. In *Proceedings of the 7th International Workshop on Spoken Language Translation, IWSLT 2010*, pages 283–289, Paris, France.
- Christian Hardmeier and Liane Guillou. 2016. A graphical pronoun analysis tool for the protest pronoun evaluation test suite. *Baltic Journal of Modern Computing*, (2):318–330.
- Christian Hardmeier, Preslav Nakov, Sara Stymne, Jörg Tiedemann, Yannick Versley, and Mauro Cettolo. 2015. Pronoun-focused MT and cross-lingual pronoun prediction: Findings of the 2015 DiscoMT shared task on pronoun translation. In *Proceedings of the Second Workshop on Discourse in Machine Translation, DiscoMT 2015*, pages 1–16, Lisbon, Portugal.
- Christian Hardmeier, Jörg Tiedemann, Preslav Nakov, Sara Stymne, and Yannick Versley. 2016. [DiscoMT 2015 Shared Task on Pronoun Translation](#). LINDAT/CLARIN digital library at Institute of Formal and Applied Linguistics, Charles University in Prague.
- Sébastien Jean, Kyunghyun Cho, Roland Memisevic, and Yoshua Bengio. 2014. On using very large target vocabulary for neural machine translation. *ArXiv e-prints*, 1412.2007.
- Sharid Loáiciga and Eric Wehrli. 2015. Rule-based pronominal anaphora treatment for machine translation. In *Proceedings of the Second Workshop on Discourse in Machine Translation*, pages 86–93, Lisbon, Portugal. Association for Computational Linguistics.
- Sharid Loáiciga. 2017. *Pronominal anaphora and verbal tenses in machine translation*. Ph.D. thesis, Université de Genève.
- Ngoc Quang Luong, Lesly Miculicich Werlen, and Andrei Popescu-Belis. 2015. Pronoun translation and prediction with or without coreference links. In *Proceedings of the Second Workshop on Discourse in Machine Translation*, pages 94–100, Lisbon, Portugal. Association for Computational Linguistics.
- Lesly Miculicich Werlen and Andrei Popescu-Belis. 2017. Validation of an automatic metric for the accuracy of pronoun translation (APT). In *Proceedings of the Third Workshop on Discourse in Machine Translation (DiscoMT)*. Association for Computational Linguistics (ACL).
- Ruslan Mitkov and Catalina Barbu. 2003. Using bilingual corpora to improve pronoun resolution. *Languages in Contrast*, 4(2):201–211.
- 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). ACL.
- Jörg Tiedemann. 2015. Baseline models for pronoun prediction and pronoun-aware translation. In *Proceedings of the Second Workshop on Discourse in Machine Translation*, pages 108–114, Lisbon, Portugal. 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*. Association for Computational Linguistics.
- Evan J. Williams. 1959. *Regression Analysis*, volume 14. Wiley, New York.