

The Two Shades of Dubbing in Neural Machine Translation

Alina Karakanta^{1,2}, Supratik Bhattacharya³, Shravan Nayak⁴,
Timo Baumann⁵, Matteo Negri¹, Marco Turchi¹

¹Fondazione Bruno Kessler, Trento, Italy ²University of Trento, Italy

³Birla Institute of Technology and Science, Pilani, India

⁴Indian Institute of Technology (BHU) Varanasi, India

⁵Universität Hamburg, Germany

{karakanta, negri, turchi}@fbk.eu,
f20170745@pulani-bits-pilani.ac.in,
pshravan.nayak.ece17@itbhu.ac.in,
baumann@informatik.uni-hamburg.de

Abstract

Dubbing has two shades; synchronisation constraints are applied only when the actor’s mouth is visible on screen, while the translation is unconstrained for off-screen dubbing. Consequently, different synchronisation requirements, and therefore translation strategies, are applied depending on the type of dubbing. In this work, we manually annotate an existing dubbing corpus (Heroes) for this dichotomy. We show that, even though we did not observe distinctive features between on- and off-screen dubbing at the textual level, on-screen dubbing is more difficult for MT (-4 BLEU points). Moreover, synchronisation constraints dramatically decrease translation quality for off-screen dubbing. We conclude that, distinguishing between on-screen and off-screen dubbing is necessary for determining successful strategies for dubbing-customised Machine Translation.

1 Introduction

Dubbing is a form of audiovisual translation (AVT) which consists in replacing the original audio track of a film dialogue with another track containing the dialogue translated in a different language. It is the preferred type of AVT in countries with large film and streaming markets, e.g. Germany, Italy, Spain and parts of Latin America (Bogucki and Díaz-Cintas, 2020). Dubbing, as a form of translation, is among the few where Machine Translation has found no steady ground yet. The first reason is the particularity of dubbing as a genre. Translation for dubbing is a text written to be spoken, therefore it should bear a close approximation to orality and reflect *oral unlabored dialog* (Chaume-Varela, 2006). However, dubbing is not spontaneous, impromptu speech but rather a carefully-prepared imitation of spoken language. This *prefabricated orality* (Baños-Piñero and Chaume, 2009) is what confers dubbing particular linguistic characteristics.

The second challenge lies in the constraints affecting the translation. The most distinctive characteristic of dubbing is the need for synchronisation when an actor’s mouth appears on screen. Isochrony requires that the duration of the source and target utterance is equal, in order for the translated dialogue to exactly fit the time during which the actor speaks. At a second level, lip-sync consists in adapting the translation to match the articulatory mouth movements of the actor, mainly matching open vowels and bilabials.¹ For example, the sentence “*I get strong off other people’s fear*” is translated as “*El miedo de los demás me da fuerzas*”, instead of a more literal translation “*Me fortalezco con el miedo de otras personas*”, in order to match the duration of the source utterance and the overlap of *f* between *fear-fuerzas*. From the above, it becomes evident that there is a clear dichotomy in dubbing strategies when an actor’s

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

¹A third type of synchronisation is *kinesic isochrony* (synchronisation with body movements) which would require a multi-modal analysis of the actions of the characters on screen and is therefore outside the scope of this paper.

mouth is visible on screen or not (off screen). While on-screen dubbing is bound by synchronisation constraints, off-screen dubbing should simply be representative of orality.

Very few works have attempted to automatise dubbing through Neural Machine Translation (NMT). Saboo and Baumann (2019) attempt to integrate the constraint of isochrony by selecting the translation which has a similar number of syllables to the source. Öktem et al. (2019) use the NMT attention mechanism to segment the translation into prosodic phrases in order to improve a Text-to-Speech system for dubbing. Federico et al. (2020) adapt an NMT system to generate translations of the same length as the source, although in terms of characters, which does not necessarily reflect duration of utterance. However, none of these works has taken into consideration the on/off-screen dichotomy.

In this work, we address, to our knowledge for the first time, the dichotomy between on-screen and off-screen dubbing in NMT. As a first contribution, we manually annotate the Heroes corpus (Öktem et al., 2018) by explicitly marking whether the actor’s mouth appears on screen or not. This annotation² is the first of its kind and will provide an invaluable resource for the study of dubbing and a benchmark for systems’ evaluation. Second, we show that synchronisation is hardly discernible at the textual level, since distinctive text features to tell apart on/off-screen dubbing are particularly hard to find. Despite this, we demonstrate that an NMT system tuned to the dubbing domain achieves significantly worse results for on-screen compared to off-screen dubbing, and that applying isochrony constraints to off-screen dubbing is detrimental for NMT quality. Our findings suggest that using a single NMT system for all “shades” of dubbing is not optimal, but viable solutions for dubbing-customised NMT should take into account the on/off dichotomy.

2 Data, Annotation and Analysis

Data: The Heroes corpus (Öktem et al., 2018) is the only freely available dubbing corpus and is based on the drama television series Heroes. It contains 7,000 single-speaker English utterance segments and their translations into Spanish, as well as time-alignments that allow for re-aligning the text with the video. Before starting the annotation, we verified that synchronisation is actually present in Heroes by watching the dubbed videos. We found that dubbing is convincing and that synchronisation is indeed present, with a stronger emphasis on preserving isochrony.

Class	Utt.s	Tokens-en	Tokens-es
ON	4,163	29,338	25,159
OFF	1,200	8,238	7,293
MIXED	1,614	18,792	16,086
MIXED-on		10,976	-
MIXED-off		7,816	-
All	6,977	56,368	48,538

Table 1: Corpus statistics in utterances and source/target tokens for the three classes.

Annotation: A first annotator watched all videos and annotated the source side of the corpus on the word-level, specifying whether speech is on-screen (mouth visible on screen) or off. Consequently, each utterance can be fully on-screen, fully off-screen, or mixed (if the actor’s mouth is visible only for some words). To validate the reliability of the annotation, a second annotator worked on 700 randomly selected utterances (10 % of the corpus). This yielded an inter-annotator agreement (Cohen’s Kappa) of $\kappa = .73$ on the word level and $.73$ on the utterance level.³ We conclude that annotation as on/off/mixed is possible with a substantial agreement across annotators. During annotation, we fixed several transcription errors and identified non-English source utterances, resulting in a total of 6,977 segments. Table 1 presents the annotation statistics.⁴ As can be seen, the majority of tokens appears on-screen (71 %; for utterances: 60 %). This imbalance is to be expected given the nature of drama television series, where the action focuses on dialogues among characters.

²The annotation is publicly available at <https://ict.fbk.eu/heroes-on-off/>

³Some deviation is expected, given that words can be very short and audio-video alignment is never perfect.

⁴Since annotation is performed only on the source, it is not possible to report Spanish tokens for the mixed category.

Ratio	ON	OFF	MIXED
Character	1.10 (0.27)	1.10 (0.38)	1.13 (0.29)
Syllable	0.90 (0.26)	0.90 (0.36)	0.93 (0.27)
Vowel	0.92 (0.28)	0.93 (0.38)	0.94 (0.29)
Consonant	1.27 (0.41)	1.28 (0.56)	1.27 (0.39)
Duration	1.04 (0.31)	1.07 (0.36)	1.04 (0.24)

Table 2: Mean and standard deviation on the source/target ratio of different units (characters, syllables, vowels vs. consonants and utterance duration in seconds).

classifier input	words	utt.s
English only	.549	.640
Spanish only		.641
English+Spanish		.634

Table 3: Area under ROC curve results for classifying words as on/off (only for English input), or utterances as on/off/mixed (for all inputs).

Analysis: We investigate characteristics that may differentiate on-screen from off-screen dubbing, focusing here only at the textual level (Table 2). We first look at relative length of translations, expressed by character and syllable ratios between source and target. We find very little differences between the mean of these ratios among the ON/OFF/MIXED classes. The same holds for the ratio between consonants and vowels, as well as for the duration of the utterances. However, a larger standard deviation of ratios for OFF suggests that there is more freedom both in terms of length and word choice when no synchronisation is needed. In order to identify possible latent characteristics (such as character sequence patterns), we also built a character-based neural sequence classifier that determines the class based on English, Spanish, or both texts.⁵ Given that classes are unbalanced, we choose Area Under ROC curve as performance metric: an ideal classifier would cover an area of 1.0, random classification yields .5. Although we find some limited success in identifying mixed utterances, our overall classification results are rather poor, as shown in Table 3. We therefore conclude that, in this dataset, there are no clear distinctive text-level features (neither direct, nor latent features that could be used by a neural network classifier) that differentiate on-screen from off-screen dubbing, i.e., that possible traces of synchronisation are hardly discernible in the textual form.

3 Neural Machine Translation for On- and Off-screen Dubbing

We now work on the hypothesis that, if there is no difference between on/off-screen dubbing, as suggested from the ratios and classification results in Section 2, an MT system should achieve comparable performance across the different classes. If this is not the case, there should be different dubbing strategies mirrored in the translation, possibly due to synchronisation constraints. To validate this hypothesis, we test the performance of an NMT model on on/off-screen dubbing.

Data: We create a test set by randomly selecting 400 sentences from each class, paying attention to avoid data leakage in the training set. From the remaining sentences, 10% from each class is sampled to form the development set. Since the number of sentences among the classes is unbalanced, we over-sample the two smaller classes (OFF and MIXED) until they reach the size of the ON class, resulting in a training corpus of 9,500 sentences. In order to obtain reliable results, we perform three rounds of cross-validation and report the mean and standard deviation of BLEU scores.

Base model: We pretrain an NMT model on a total of 67M parallel sentences from the OPUS project,⁶ containing OpenSubtitles (Lison and Tiedemann, 2016), Europarl (Koehn, 2005), GlobalVoices⁷, MuST-C (Di Gangi et al., 2019), WIT³ (Cettolo et al., 2012). All data are segmented into subword units using SentencePiece (Kudo and Richardson, 2018) with a 40K joint vocabulary. The model is based on the Transformer (big) architecture (Vaswani et al., 2017) of the fairseq toolkit (Ott et al., 2019). It is trained with label smoothed cross-entropy and 0.1 label smoothing. For optimisation, we use Adam (Kingma and Ba, 2015) with an initial learning rate of 1×10^{-7} , which increases linearly up to 0.005 for 4000 warm-up steps, and then decreases with the inverse square root of the training step. Dropout is set to 0.3

⁵Implementation details as well as the code can be found at <https://github.com/timobaumann/visualHeroesAnnotations/tree/master/text-classifier>

⁶<http://opus.nlpl.eu/>

⁷<http://opus.nlpl.eu/GlobalVoices-v2017q3.php>

for all layers except for the attention layer, where dropout is set to 0.1. The base model achieves a BLEU score of 31.47 on the WMT’13 English→Spanish test set (Bojar et al., 2013).

Proposed models: We customise the base NMT model to the dubbing domain by fine-tuning on the Heroes data. Given the scarcity of training material, we concatenate data from all classes and follow two different strategies: 1) simply fine-tuning on all the dubbing data (*ft-heroes*) and 2) fine-tuning with *target forcing*, as in multilingual translation (Johnson et al., 2017). This technique has been shown to be beneficial not only for mixing data from different languages, but also to achieve style transfer (Niu et al., 2017) and to control attributes of the target text, such as politeness (Sennrich et al., 2016). Since the third class (MIXED) is essentially a mix of ON and OFF, we experiment with two tagging strategies for target forcing: a) we distinguish the 3 classes with the tags <ON>, <OFF> and <MIX> (*ft-3tags*), and b) we use only the <ON> and <OFF> tags, by annotating the ranges for on/off-screen text in the MIXED class (*ft-2tags*). The second strategy is inspired by the use of dubbing symbols in the actual dialogue lists (Chaume, 2012), therefore it explores the possibility of a direct application in the dubbing industry.

Additionally, we apply the previous work attempting to integrate dubbing constraints in NMT (Saboo and Baumann, 2019) to our *ft-2tags* model (*syll-rescore*). We use a beam search of 20 and select the hypothesis with the closest source/target syllable ratio⁸ to the target syllable ratio observed in our data (0.9). Compared to previous work, we use a smaller beam (20 instead of 50) because we found that, in 90% of the sentences, at least one candidate matching the target ratio was found among the hypotheses.

4 Results

The BLEU scores for the models described above are shown in Table 4.

On/Off-screen dubbing: Fine-tuning on the Heroes data (*ft-heroes*) gives an increase from 21.30 to 25.89 BLEU for ON (+19%), from 20.07 to 24.95 for MIXED (+19%) and from 24.04 to 30.40 for OFF (+24%). The score for the OFF class is higher under all conditions, showing that off-screen dubbing is less challenging for NMT. Fine-tuning on in-domain data is in general expected to bring performance gains. Based on the evidence shown, we cannot establish whether the model is simply fine-tuned to the style and vocabulary of the Heroes TV series, or to the dubbing task. Hypothetically, if we could fine-tune on another dubbing dataset, we would still likely see performance gains. This would be an indication that, even though the *base* model is trained on corpora of the spoken domain, dubbing is a particular sub-domain with its own particularities in terms of style and orality.

Tagging: Target forcing gives a ~ 1 BLEU point increase, showing that it helps the NMT system adapt to different styles and particularities even inside this sub-domain. For the ON class, the increase comes only when the MIXED data are separated in ON and OFF (*ft-2tags*), since this is the only condition when NMT receives more ON data. In this condition, there is a slight drop in BLEU for the MIXED class, possibly because the mix of tags increases the complexity of the source sentence.

Syllabic rescoring: When applying isochrony constraints to select the translation candidate with the most similar number of syllables to the source, the translation performance is compromised. This drop is in line with Saboo and Baumann (2019) who attempted to balance the trade-off between isochrony and translation quality, but without making any on/off-screen distinction. Our results show an interesting tendency; while the drop for ON and MIXED when applying synchronisation constraints is 3 BLEU points, for OFF the score drops by 7 points. As a result, the translation performance for all the classes is flattened to ~ 23 -24 BLEU points. This finding suggests that for off-screen dubbing, applying synchronisation constraints is not required but rather detrimental. On the other hand, the compromise for the classes containing on-screen dubbing is smaller, which advocates for the presence of some isochrony constraint for ON and MIXED. Still, the drop in translation performance indicates that isochrony is not the only constraint affecting the translation, but the need for adaptation to the articulatory movements (lip-sync) at the level of phonemes should be further explored in NMT for dubbing.

⁸Computed with the Pyphen package: <https://pyphen.org/>

Model	Total	ON	OFF	MIXED
base	21.32 (0.61)	21.30 (0.5)	24.04 (0.44)	20.07 (1.22)
ft-heroes	26.54 (3.09)	25.89 (2.6)	30.40 (4.2)	24.95 (3.09)
ft-3tags	27.45 (3.1)	25.82 (2.23)	31.11 (4.01)	27.01 (3.01)
ft-2tags	27.50 (3.09)	26.63 (2.47)	30.84 (4.3)	26.33 (3.76)
syll-rescore	22.80	23.31	23.71	24.11

Table 4: BLEU scores for different systems for the entire test set (Total) and the three classes. Mean and StDev (in parentheses) of 3 cross-validation rounds.

Model	ON	OFF	MIXED
Human	347	363	398
ft-2tags	304	339	348
syll-rescore	335	370	374

Table 5: Perplexity on a general Spanish language model for human dubbing and two MT outputs.

5 Analysis

The BLEU scores and the improvements from fine-tuning are higher for the OFF class. Our first hypothesis was that this difference is due to different contents/styles in the on/off test sets; for example, off-screen contains more narratives that are easier to translate, whereas on-screen more informal language. However, the majority of segments come from dialogues between characters, and narratives account only for around 10 segments in the whole corpus. Since the styles of the source segments are comparable, the lower BLEU scores for the ON and MIXED class suggest that (partially and fully) on-screen dubbing is highly constrained translation, which depends not only on semantics, but on factors to which the NMT model does not have access, such as phonetics and visemics. Therefore, this difference should be investigated in the interplay between the source and target in terms of translation solutions.

To account for the difference in BLEU, we compute perplexity of the human dubbing and NMT outputs on a 5-gram language model (Heafield et al., 2013) trained on a corpus of general Spanish language (Cardellino, 2019) (Table 5). A lower perplexity and lower BLEU score for ON suggests that there is increased naturalness for on-screen human dubbing, imposed by the need to create realistic dialogues, in line with the factors above. This results in creative translations that rely less on the lexical choice, structure and word order of the source text; NMT systems normally lack this creativity. Indeed, for human dubbing the conformity to the target language norms has been claimed to be more important than following the source text structure or generating a perfect lip-sync (Chaume, 2012; Pavesi, 2008). On the other hand, a higher BLEU and higher perplexity for OFF indicate less creative solutions, closer to the source, and therefore more predictable for the NMT system. With a low BLEU and high perplexity, MIXED seems to be the most particular class, where the mix of dubbing strategies creates a linguistic hybrid.⁹ This tendency is copied in the NMT outputs, which however are less surprising under all conditions, suggesting that they are more plausible as general language, but they do not reflect the orality of dubbed texts. The high perplexity for OFF under *syll-rescore* is another indication that syllabic rescoring is harmful, since it leads to translations that are erroneous or not plausible in the target language.

6 Conclusions and future work

We have explored for the first time the characteristics of on/off-screen dubbing from a computational perspective. To this end, we annotated the Heroes corpus, distinguishing between the two shades of dubbing. Our annotation is an important starting point towards fulfilling the multimodal requirements necessary for fully functional dubbing engines which will incorporate detection of visuals for identifying the shade and deciding whether to apply synchronisation (Nayak et al., 2020). Despite the lack of distinctive features at the textual level in our dataset, NMT seems to suffer from a (still) elusive difference between on/off-screen dubbing, as witnessed by lower BLEU scores for ON and MIXED. We have further shown that isochrony constraints significantly hurt NMT performance for off-screen dubbing. Understanding the language of dubbing is still an open problem. We hope that our findings will set a first cornerstone towards the successful integration of NMT in dubbing. In the future, we will investigate isochrony and especially lip-sync from a phonetic perspective.

⁹On a language model of English, no difference in perplexity was found between ON and OFF, (0.1 points) while perplexity was higher for MIXED (10 points).

References

- Rocío Baños-Piñero and Frederic Chaume. 2009. Prefabricated orality a challenge in audiovisual translation. In G. Nadiani M. Giorgio Marrano and C. Rundle, editors, inTRAlinea Special Issue: The Translation of Dialects in Multimedia. inTRAlinea. Last accessed: June 2020.
- Łukasz Bogucki and Jorge Díaz-Cintas. 2020. An excursus on audiovisual translation. In Ł. Bogucki and M. Deckert, editors, The Palgrave Handbook of Audiovisual Translation and Media Accessibility, pages 11–32. Palgrave Studies in Translating and Interpreting.
- Ondřej Bojar, Christian Buck, Chris Callison-Burch, Christian Federmann, Barry Haddow, Philipp Koehn, Christof Monz, Matt Post, Radu Soricut, and Lucia Specia. 2013. Findings of the 2013 Workshop on Statistical Machine Translation. In Proceedings of the Eighth Workshop on Statistical Machine Translation, pages 1–44, Sofia, Bulgaria. Association for Computational Linguistics.
- Cristian Cardellino. 2019. Spanish Billion Words Corpus and Embeddings. Last accessed: June 2020.
- Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. WIT³: Web Inventory of Transcribed and Translated Talks. In Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT), pages 261–268, Trento, Italy, May.
- Frederic Chaume-Varela. 2006. Dubbing. In Keith Brown, editor, Encyclopedia of Language & Linguistics, pages 6 – 9. Elsevier, Oxford, Second edition.
- Frederic Chaume. 2012. Audiovisual Translation: Dubbing. Translation practices explained. St. Jerome Pub.
- Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. MuST-C: a Multilingual Speech Translation Corpus. 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 2012–2017, Minneapolis, Minnesota. Association for Computational Linguistics.
- Marcello Federico, Robert Enyedi, Roberto Barra-Chicote, Ritwik Giri, Umut Isik, Arvindh Krishnaswamy, and Hassan Sawaf. 2020. From speech-to-speech translation to automatic dubbing. In Proceedings of the 17th International Conference on Spoken Language Translation, pages 257–264, Online. Association for Computational Linguistics.
- Kenneth Heafield, Ivan Pouzyrevsky, Jonathan H. Clark, and Philipp Koehn. 2013. Scalable modified Kneser-Ney language model estimation. In Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pages 690–696, Sofia, Bulgaria. Association for Computational Linguistics.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. Google’s Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation. Transactions of the Association for Computational Linguistics, 5:339–351.
- Diederik Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In 3rd International Conference for Learning Representations.
- Philipp Koehn. 2005. Europarl: a parallel corpus for statistical machine translation. In Proceedings of the Tenth Machine Translation Summit, pages 79–86. Asia-Pacific Association for Machine Translation.
- 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, November. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16), pages 923–929, Portorož, Slovenia, May. European Language Resources Association (ELRA).
- Shravan Nayak, Timo Baumann, Supratik Bhattacharya, Alina Karakanta, Matteo Negri, and Marco Turchi. 2020. See me speaking? differentiating on whether words are spoken on screen or off to optimize machine dubbing. In Companion Publication of the 2020 International Conference on Multimodal Interaction (ICMI ’20 Companion), page 5. ACM, New York, NY, USA, October.

- Xing Niu, Marianna Martindale, and Marine Carpuat. 2017. A study of style in machine translation: Controlling the formality of machine translation output. In Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pages 2814–2819, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Alp Öktem, Mireia Farrús, and Antonio Bonafonte. 2018. Bilingual Prosodic Dataset Compilation for Spoken Language Translation. In Proceedings of IberSPEECH 2018, pages 20–24.
- Alp Öktem, Mireia Farrús, and Antonio Bonafonte. 2019. Prosodic Phrase Alignment for Machine Dubbing. In Proceedings of INTERSPEECH 2019, pages 4215–4219.
- 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 NAACL-HLT 2019: Demonstrations.
- Maria Pavesi. 2008. Spoken language in film dubbing: Target language norms, interference and translational routines. In Delia Chiaro, Christine Heiss, and Chiara Bucaria, editors, Between Text and Image: Updating research in screen translation, number 78, pages 79–99. Benjamins Translation Library.
- Ashutosh Saboo and Timo Baumann. 2019. Integration of dubbing constraints into machine translation. In Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers), pages 94–101, Florence, Italy. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Controlling politeness in neural machine translation via side constraints. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 35–40. 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, pages 6000–6010.