

Inference Strategies for Machine Translation with Conditional Masking

Julia Kreutzer and George Foster and Colin Cherry

Google Research

{jkreutzer, fosterg, colincherry}@google.com

Abstract

Conditional masked language model (CMLM) training has proven successful for non-autoregressive and semi-autoregressive sequence generation tasks, such as machine translation. Given a trained CMLM, however, it is not clear what the best inference strategy is. We formulate masked inference as a factorization of conditional probabilities of partial sequences, show that this does not harm performance, and investigate a number of simple heuristics motivated by this perspective. We identify a thresholding strategy that has advantages over the standard “mask-predict” algorithm, and provide analyses of its behavior on machine translation tasks.

1 Introduction

The widely successful masked language modeling paradigm popularized by BERT (Devlin et al., 2019) has recently been adapted to conditional masked language model (CMLM) training for semi-autoregressive sequence generation (Ghazvininejad et al., 2019), where model predictions are conditioned on the complete input sequence and the observed (non-masked) portion of the output sequence. The CMLM’s simplicity and its clear links to the very active field of linguistic representation learning are advantages over its semi-autoregressive competitors, such as iterative refinement of token sequences (Lee et al., 2018), refinement of non-linguistic intermediate representations (Kaiser et al., 2018; Shu et al., 2020) and learning to predict parallel edit operations (Stern et al., 2019; Gu et al., 2019).

It is not obvious how to best perform inference with the CMLM. Starting from a partially-observed output sequence, the optimal choice to complete it within a single step would be to generate the most likely token at each unobserved (masked) position independently. However, it is less clear how

to progress from an initial, completely masked sequence to a final hypothesis semi-autoregressively over a number of steps, with each successive step unmasking new context for the next. This requires not only ordering the tokens for generation, but also making decisions about how many tokens to simultaneously predict in each step.

Ghazvininejad et al. (2019) propose the *mask-predict* algorithm, which iteratively generates fresh model predictions for all masked positions, and then unmask a predefined number of the most likely predictions. Given a fixed number of iterations, a decaying schedule determines how many predictions to unmask in each iteration. Each successive iteration provides mode-breaking (Gu et al., 2018) context for the next. By fixing the number of iterations, this approach allows for constant-time semi-autoregressive decoding.

The fixed-iteration strategy is very practical and has yielded empirical success in a range of machine translation experiments, but there is no guarantee that it is optimal. The tokens to be unmasked on a given iteration are all predicted independently, and therefore might contain repeated words, or words with low model confidence. These issues can be mitigated by later re-masking a token to repair it (Ghazvininejad et al., 2019) or by adapting the model to incorrect contexts (Ghazvininejad et al., 2020).

We instead adopt a fully probabilistic view of the masked prediction sequence, which we enable by simply disallowing the re-masking of previously unmasked tokens. This view guides us to a heuristic inference schedule that selects sets of unmasked tokens according to a threshold on the product of their conditionally independent model probabilities. This heuristic naturally slows down in the situations mentioned above, and speeds up in the presence of high confidence, which allows us to achieve favorable quality-to-speed trade-offs.

We focus on strengthening the CMLM inference (Section 3) while leaving its training algorithm unchanged (Section 2), and maintaining much of the structure of the original inference strategy. For our experiments on machine translation (Section 4), we compare inference heuristics in terms of their quality-speed trade-offs. We analyze the development of quality over iterations, and the influence of sentence length. With examples of unmasking schedules we furthermore illustrate the role of mode breaking through choosing the right contexts.

2 CMLM Model and Training

The CMLM is a model for $p(Y_{mask}|Y_{obs}, X)$, the probability of masked tokens Y_{mask} given a partially observed output sequence Y_{obs} and an input sequence X . Y_{mask} and Y_{obs} are sets of tokens at specified positions that together form a complete output sequence Y : $Y_{mask} = Y \setminus Y_{obs}$. The model is implicitly conditioned on output sequence length $N = |Y|$, and the tokens in Y_{mask} are conditionally independent: $p(Y_{mask}|Y_{obs}, X) = \prod_{y_i \in Y_{mask}} p(y_i|Y_{obs}, X, N)$. During training, masks are placed randomly: First, the mask size $S \in \{1, \dots, N\}$ is sampled from a uniform distribution, then S positions are randomly chosen to define the subsets Y_{obs} and Y_{mask} . Cross-entropy loss is incurred via $p(y_i|Y_{obs}, X)$ for each $y_i \in Y_{mask}$. An additional classifier on top of encoder representations is trained to predict the output length N .

3 CMLM Inference

Inference starts with a context of only MASK tokens. Until a stop condition is met, decoder predictions iteratively replace a subset of these in selected positions (“unmasking”). With a single iteration, inference is non-autoregressive; when the number of iterations T is less than the sentence length N it is semi-autoregressive; and when $N = T$ it is fully autoregressive. Due to the use of a uniform distribution over reference contexts, training is agnostic to these different regimes.

In general, we seek to minimize T without trading off too much quality. The challenge in doing so is to identify the subset of predictions that are most likely to provide suitable conditioning context for future iterations (Mansimov et al., 2019). Structural or linguistic dependencies in the output may also play an important role for resolving linguistic ambiguities (Martins and Kreutzer, 2017).

t	$M^{(t)}$	$Y^{(t)}$	$p(Y^{(t)} Y^{(<t)}, M^{(\leq t)}, X)$
0	{1,2,3}	{}	–
1	{}	{a,b,c}	p(alX) p(blX) p(clX)
0	{1,2,3}	{}	–
1	{2,3}	{a}	p(alX)
2	{3}	{b}	p(bla,X)
3	{}	{c}	p(cla,b,X)
0	{1,2,3}	{}	–
1	{2}	{a,c}	p(alX) p(clX)
2	{}	{b}	p(bla,c,X)

Figure 1: Computations for $p(Y, M|X)$, where $Y = \{a, b, c\}$, for various masking sequences M . The first sequence is fully non-autoregressive, and the second is the standard left-to-right autoregressive factorization.

For example, in German it might be harder to first generate the determiner before knowing the grammatical gender of the head word (see examples in Figure 6).

The length predictor first predicts b different lengths, then one hypothesis is decoded for each length independently using the iterative process just outlined. The hypothesis with the highest length-normalized model score is selected as the output. We refer to b as the *length beam* in the following.

3.1 Update Strategies

The CMLM can make predictions at all positions, whether they correspond to masked input or not. This lends itself to various strategies for choosing how to update current predictions and masks:¹

- *update-all*: update tokens and scores at all positions, no constraint on new mask²
- *update-masked*: update tokens at masked positions only, no constraint on new mask³
- *update-masked-sub*: update tokens at masked positions only, new mask must be a subset of the current one

In this paper we focus on the *update-masked-sub strategy*. It is empirically competitive (Section 4.1), and interesting because it corresponds to a valid probabilistic factorization of the target sequence, governed by a latent variable $M = M^{(0)} \dots M^{(T)}$

¹In all cases we assume predictions to be the most likely words at each position, and scores to be the corresponding probabilities.

²This corresponds to a masked version of iterative refinement (Lee et al., 2018).

³This is the strategy used by Ghazvininejad et al. (2019).

which represents the sequence of masking decisions:

$$p(Y, M|X) = \prod_{t=1}^T p(Y^{(t)}|Y^{(<t)}, M^{(\leq t)}, X) \times p(M^{(t)}|Y^{(<t)}, M^{(<t)}, X), \quad (1)$$

where $M^{(0)} = \{1, \dots, N\}$, $M^{(T)} = \{\}$, $M^{(t)} \subset M^{(t-1)}$, and $Y^{(t)}$ is the set of tokens *unmasked* on the t th iteration. Figure 1 illustrates this computation for various choices of M .⁴

The class of inference strategies we explore can thus be seen as *greedy search* for the mostly likely factorization, subject to a constraint on the number of iterations: at each iteration, we choose a subset of tokens to add to the current hypothesis, balancing high model probabilities with the risk of making an error and degrading future predictions. Because tokens are predicted independently, the risk of an error grows with the size of the subset.

3.2 Unmasking Heuristics

Under the *update-masked-sub* constraint, the role of greedy inference heuristics is to choose which positions to unmask, given a full set of predictions for all currently-masked positions. The *mask-predict* strategy of Ghazvininejad et al. (2019) chooses the $\lceil N/T \rceil$ highest-probability tokens, in order to finish in a constant T iterations, regardless of N . This generates more tokens per iteration for long sentences, which may not be ideal for sentences with complex structure. To measure its effect, we propose a variant that unmask a constant K tokens per iteration, in order to achieve approximately K -fold speedup over autoregressive performance, independent of hypothesis length.

Unmasking highest-ranked tokens according to probability is reasonable, but it ignores the magnitude of the probabilities, creating the potential for selecting tokens in which the model has low confidence, and vice versa. To address this, we design several simple thresholding strategies that vary the number of tokens per iteration, ideally generating more when the conditioning context licences many confident predictions, and fewer otherwise.

1. The most straightforward strategy, *thresh*, un-masks all tokens with probabilities greater than a given threshold τ .

⁴Note that a probabilistic interpretation enables an unconstrained search for the most probable output, or for the unmasking sequence that assigns highest probability to a reference output, options we do not pursue in this paper.

2. The *comb-thresh* strategy unmask the largest set of highest-ranked tokens Y whose joint probability $p(Y) > \tau$.
3. Finally, in order to account for lower-ranked predictions, the *fcomb-thresh* strategy un-masks the largest set Y for which $p(Y) * (1 - p(\bar{Y})) > \tau$, where Y consists of the highest-ranked tokens, and \bar{Y} is its complement.

All threshold strategies unmask the single highest-ranked token in contexts where the threshold criterion is not met.

4 Experiments

Our CMLM is implemented with a base Transformer (Vaswani et al., 2017) built on a TensorFlow implementation of (Ghazvininejad et al., 2019). The input to the decoder is Y_{obs} , with MASK tokens at masked positions, and the output is Y_{mask} , predictions for all masked positions without future attention masking. We use data from WMT14 en \leftrightarrow de (Bojar et al., 2014) and WMT17 zh \leftrightarrow en (Bojar et al., 2017) with a sentence piece vocabulary of 32k, focusing mainly on en \rightarrow de, and providing results for all pairs in appendix A. The CMLM is trained on distilled training data from an autoregressive Transformer and initialized with its parameters.

4.1 Update Strategies

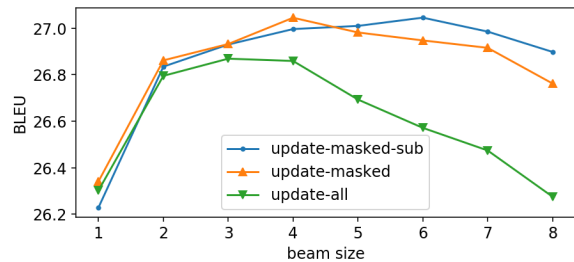


Figure 2: Performance of update strategies

Figure 2 shows the performance of the update strategies described in section 3.1 versus length beam b . All strategies use the mask-predict heuristic with a fixed 10-iteration limit. As beam size increases past 2, the update-masked strategies increasingly dominate, indicating that their scores are more reliable for choosing among length hypotheses. There is no significant difference between the two variants of update-masked. This suggests that our probabilistic factorization constraint (*update-masked-sub*) does not hurt in practice.

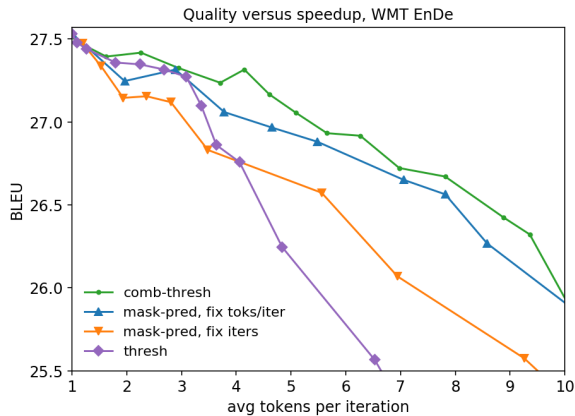


Figure 3: Heuristics with length-beam=5. The top border of the graph represents the performance of an autoregressive Transformer with beam=5. Different points in the graph correspond to different hyperparameter settings, varying the hyperparameter which controls the speed for each heuristic (T for fixed-iteration *mask-predict*, K for variable-iteration *mask-predict*, and τ for thresholding heuristics).

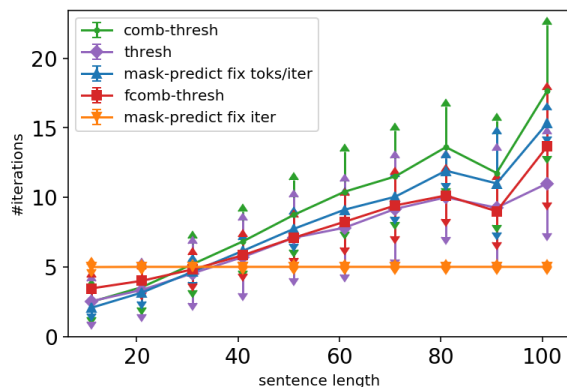


Figure 4: Number of iterations on the test set in relation to (oracle) sentence length (tokens) when generating on average ca. 5 tokens per iteration. Error bars indicate the standard deviation.

4.2 Heuristics

To compare the speed-quality trade-off of different heuristics on an equal footing, we vary the values of the hyper-parameter that controls speedup: T for fixed-iteration *mask-predict*, K for variable-iteration *mask-predict*, and τ for thresholding strategies. In each case, we measure the resulting speedup as the total number of tokens in the test set divided by the total number of iterations required for all sentences,⁵ and corpus BLEU on the output of the last iteration.

⁵This is theoretical speedup, and we make no claims that it can be attained in practice, an objective that would likely require significant engineering effort.

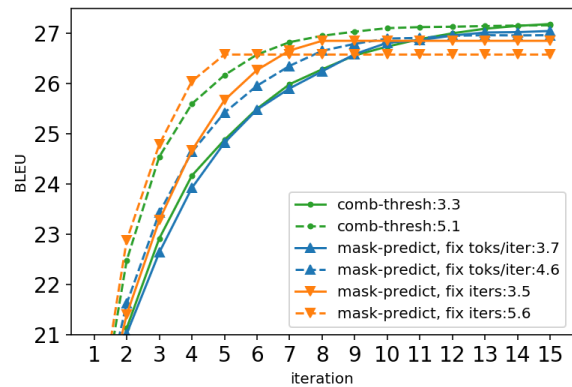


Figure 5: Development of BLEU over iterations comparing heuristics for two different generation speeds. The generation speed is expressed in average number of tokens per iteration, e.g. *comb-thresh*:3.5 stands for the *comb-thresh* heuristic with a threshold value set so that 3.5 tokens are generated per iteration on average.

Figure 3 compares heuristics using 5 length candidates. First of all, fixed- K *mask-predict* beats fixed- T by a substantial margin (especially at higher speeds), indicating that it is worth allocating more iterations for longer sentences. Second, the *comb-thresh* strategy has a small but consistent advantage over fixed- K *mask-predict* across all speeds. This strategy exhibits a roughly 4x gain while sacrificing less than 0.3 BLEU relative to the equivalent autoregressive Transformer (27.6 BLEU).

Both *thresh* and *fcomb-thresh* underperform. Despite their superficial similarity to *comb-thresh*, they perform much worse; this holds for other language pairs as well (Figure 7 in Appendix A). For *thresh*, the poor performance as speedup increases reflects many relatively low-probability tokens exceeding lower thresholds, a condition that is penalized by all other heuristics, which take rank into account. For *fcomb-thresh* the effect is more subtle; we believe that it is due to the probabilities of lower-ranked tokens having worse calibration, leading to less reliable unmasking decisions.

A practical impediment to a thresholding strategy is that it does not provide direct control over desired speedup: this must be identified by tuning τ appropriately on a development set. However, we found that dev and test speedups were well correlated across speedups ranging from 1 to 11, with the largest absolute error being 0.8 (11.1 speedup on dev versus 10.3 on test), and the average error being 0.3.




Figure 6: Unmasking over iterations for source sentence "A job like this is not something you achieve overnight," said Gerster in praise of the annual financial statement. with reference "So ein Werk schüttelt man nicht einfach aus dem Ärmel", lobte Gerster mit Blick auf die Jahresrechnung. for *mask-predict* with $T = 5$ above and *comb-thresh* with $\tau = 0.1$ below. Predicted tokens are shown when they differ from the previous iteration. Their background color indicates the model score, with yellow indicating high scores, and dark blue low scores.

4.3 Analysis

Having freed the heuristics from a globally imposed iteration limit for constant-time decoding as in the original *mask-predict* inference heuristic, we observed better quality-speed trade-offs in the above discussed results. Intuitively, we would expect the heuristics to allocate more iterations for longer sentences and save iterations on shorter sentences. Figure 4 shows how many iterations the models spend on sentences in relation to their length. For a fair comparison, the generation is constrained to oracle output lengths, and we set the hyperparameters such that they result in the same generation speed (5 tokens per iteration on average). We see that flexible-iteration strategies spend fewer iterations on sentences up to a length of around 30 when compared to a fixed-iteration strategy. *comb-thresh* spends on average the largest number of iterations on longer sentences (which pays off in terms of quality, see Figure 3), while *thresh* spends even fewer iterations on longer sentences than the *mask-predict* model.

The development of BLEU over iterations for comparable generation speeds across heuristics is shown in Figure 5.⁶ We can see that speedier generation gives a faster initial increase in translation quality over iterations in exchange for slightly lower final quality (dashed vs solid lines). *Mask-predict* levels off early after reaching its fixed number of iterations, but climbs quickly before that point due to an averaging effect over short sentences. Fixed-K *mask-predict* and *comb-thresh* both extract useful work out of each iteration, with *comb-thresh* maintaining a slight edge over all iterations, especially at higher generations speeds.

Figure 6 shows an example for generation strate-

gies under *mask-predict* and *comb-thresh* (see appendix B). They illustrate the workings of iterative decoding and main differences between strategies: Iterative decoding is crucially needed to resolve subject-verb agreement (e.g. "man [...] erreichen" (generic "you") vs. "Sie [...] erreichen" (formal "you") in ex. 2) and rough sentence structure (e.g. placement of the comma), and offers room for less literal translations ("von heute auf morgen" (literally "from today to tomorrow") rather than "über Nacht" (literally "over night") in ex. 1). The two tokens "Ger" and "ster" (a name) show how the correct conditioning changes model scores in both cases: After the former token is predicted, the probability for the latter increases drastically, since its only valid position in the sentence is there. While both strategies use the same number of iterations to generate this translation, one can see that it pays off for *comb-thresh* to unmask certain tokens earlier ("ab", "Lob"), which allows a valid resolution of neighboring tokens ("abschluss", "zum").⁷

5 Conclusion

We investigated inference strategies for machine translation based on CMLM with a focus on the trade-off between generation speed and quality. We introduce a perspective which views generation sequences as probabilistic factorizations of the final output sequence, and use it to analyze and extend previous heuristics. Our new heuristics achieve better speed/quality balance by flexibly adjusting the number of total iterations, and by taking the probabilities of sets of tokens into account. For future work we would like to explore if their success transfers to other generation tasks with MLMs where inference efficiency is a concern.

⁶Each line on this graph is produced by doing inference with a particular hyperparameter setting, and recording BLEU for the greedily predicted tokens after each iteration.

⁷A typo ("abschluss" vs "abschluss") is introduced by choosing the "ab" sub-word rather than "a", likely contributing to the model uncertainty in this area.

References

- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. [Findings of the 2017 conference on machine translation \(wmt17\)](#). In *Proceedings of the Second Conference on Machine Translation (WMT)*, Copenhagen, Denmark.
- Ondrej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation (WMT)*, Baltimore, Maryland.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, Minneapolis, Minnesota.
- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. [Mask-predict: Parallel decoding of conditional masked language models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Hong Kong, China.
- Marjan Ghazvininejad, Omer Levy, and Luke Zettlemoyer. 2020. [Semi-autoregressive training improves mask-predict decoding](#). *arXiv preprint arXiv:2001.08785*.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor O.K. Li, and Richard Socher. 2018. [Non-autoregressive neural machine translation](#). In *International Conference on Learning Representations (ICLR)*, Vancouver, Canada.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. [Levenshtein transformer](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, Vancouver, Canada.
- Lukasz Kaiser, Samy Bengio, Aurko Roy, Ashish Vaswani, Niki Parmar, Jakob Uszkoreit, and Noam Shazeer. 2018. [Fast decoding in sequence models using discrete latent variables](#). volume 80 of *Proceedings of Machine Learning Research*, pages 2390–2399, Stockholm, Sweden.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. [Deterministic non-autoregressive neural sequence modeling by iterative refinement](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium.
- Elman Mansimov, Alex Wang, and Kyunghyun Cho. 2019. [A generalized framework of sequence generation with application to undirected sequence models](#). *arXiv preprint arXiv:1905.12790*.
- André F. T. Martins and Julia Kreutzer. 2017. [Learning what’s easy: Fully differentiable neural easy-first taggers](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Copenhagen, Denmark.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation (WMT)*, Brussels, Belgium.
- Raphael Shu, Jason Lee, Hideki Nakayama, and Kyunghyun Cho. 2020. [Latent-variable non-autoregressive neural machine translation with deterministic inference using a delta posterior](#). In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, New York, New York.
- Mitchell Stern, Will Chan, Jamie Kiros, and Jakob Uszkoreit. 2019. [Insertion transformer: Flexible sequence generation via insertion operations](#). In *International Conference on Machine Learning (ICML)*, Long Beach, California.
- 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 30 (NeurIPS)*, Long Beach, California.

A Experiments on different languages

Pre-processing for all data sets follows the procedure described in (Vaswani et al., 2017).

Figure 7 shows the results for heuristics on all language pairs. As in our main experiments, quality is measured with tokenized BLEU, except for en→zh, where we use SacreBLEU (Post, 2018). In three of the language pairs, we observe a similar pattern to en→de: *comb-thresh* has a slight but consistent advantage over *mask-predict*, with the fixed tokens/iteration version of *mask-predict* doing consistently better than the fixed iteration version. On en→zh, all three methods perform similarly.

B Examples

Figure 8 provides more examples for *mask-predict* (Figure 8a) and *comb-thresh* (Figure 8b) heuristics under different hyperparameter settings, complementing the ones displayed in Figure 6. The source sentence is "A job like this is not something you achieve overnight," said Gerster in praise of the annual financial statement., and the reference "So ein Werk schüttelt man nicht einfach aus dem Ärmel", lobte Gerster mit Blick auf die Jahresrechnung.. Predicted tokens are printed out when they differ from the previous iteration. Their background color indicates the model score, with yellow indicating high scores, and dark blue low scores.

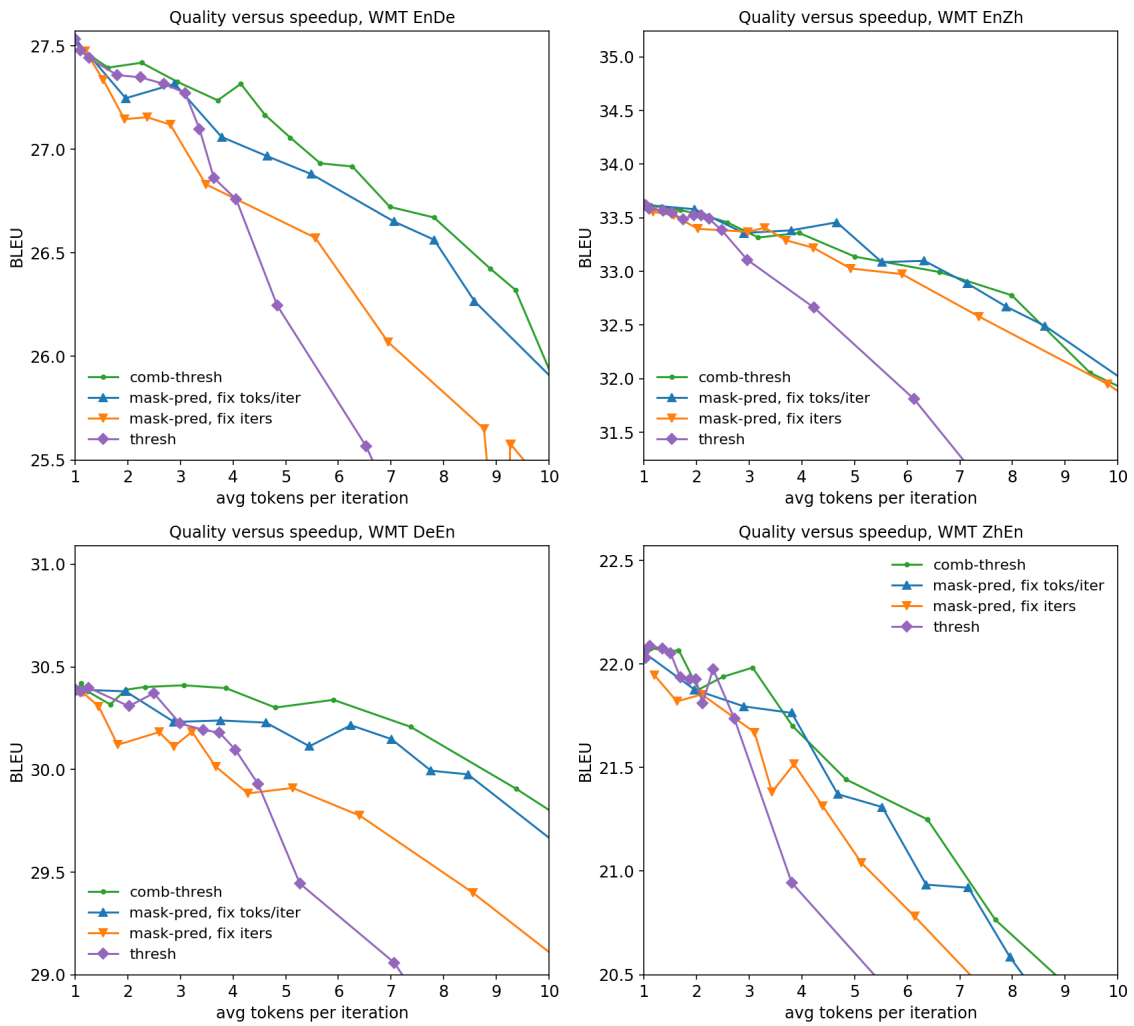
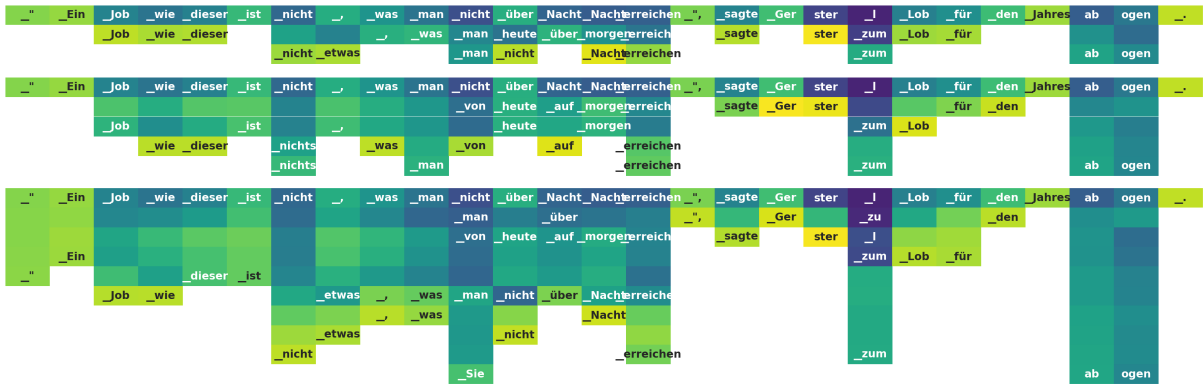



Figure 7: Results on all language pairs, with length-beam 5. The top border of each graph represents performance of the equivalent autoregressive Transformer.



(a) Configurations of *mask predict* with fixed number of iterations: (1) $T = 3$, (2) $T = 5$, (3) $T = 10$.



(b) Configurations of *comb-thresh*: (1) $\tau = 0.002$, (2) $\tau = 0.1$, (3) $\tau = 0.4$.

Figure 8: Example unmasking schedules. Predicted tokens are shown when they differ from the previous iteration. Their background color indicates the model score, with yellow indicating high scores, and dark blue low scores.