Reverse Thinking in Large Language Models

Anonymous ACL submission

Abstract

Humans are accustomed to reading and writing in a forward manner, and this natural bias extends to text understanding in auto-regressive large language models (LLMs). This paper investigates whether LLMs, like humans, struggle with reverse thinking, specifically with reversed text inputs. We found that publicly available pre-trained LLMs cannot understand such inputs. However, LLMs trained from scratch with both forward and reverse texts can understand them equally well during inference. Our case study shows that different-content texts result in different losses if input (to LLMs) in different directions—some get lower losses for forward while some for reverse. This leads us to a simple and nice solution for data selection based on the loss differences between forward and reverse directions. Using our selected data in continued pretraining can boost LLMs' performance by a large margin for the task of Massive Multitask Language Understanding.

1 Introduction

LLMs (Touvron et al., 2023; Jiang et al., 2023) have shown impressive capabilities across diverse natural language processing tasks and beyond. These capabilities are primarily attributed to the learning of extensive corpora that cover general world knowledge (Kaplan et al., 2020). These corpora are created in human society and often demonstrate human bias, including the inherently forward-oriented human cognition (Bergen and Chan, 2005; De Kerckhove and Lumsden, 2013), *e.g.*, reasons precede outcomes and solutions are deduced from given information. In contrast, reverse thinking presents more cognitive challenges due to its contradiction with innate commonsense and human logic. Prior studies (Taylor and McNemar, 1955; Gao and Wang, 2019; Sweller, 2020) indicate that reverse thinking can substantially improve cognitive abilities. This raises the question of *can LLMs do reverse thinking or will they face similar challenges as humans?*, and *could reverse thinking benefit the learning of LLMs?*.

To study this, we simulate reverse-thinking data by reversing entire paragraphs or documents at the token level. Please note that this is the simplest and extreme way but may not be the optimal way of generating reverse thinking. We train LLMs with these simulated texts and conduct a comprehensive analysis. Overall results indicate that LLMs learn forward- and reverse-thinking texts equally well when trained from scratch. However, performance varies across text samples. Some are suited to reverse thinking, while others favor forward thinking. Notably, we found that texts suited for reverse thinking are often high-quality and more logically coherent. Training on them, the original "forward-thinking" LLMs can be improved. We conduct the empirical validation on the task of Massive Multitask Language Understanding (MMLU) (Hendrycks et al., 2020). In summary, this paper has two main contributions. First, it demonstrates and analyzes the performance of LLMs learned from forward- and reverse-thinking texts. Second, it shows that the data selected based on the losses of forward- and reverse-thinking texts can further improve LLMs.

2 Related Work

In this paper we utilize the reverse text for model training. Previous work on reverse input falls into two main areas. The first area involves using reverse text in machine translation. Studies show that using decoders to process text both left-to-right and right-to-left within an encoder-decoder framework improves machine translation performance (Zhou et al., 2019; Gu et al., 2019), a finding later extended to LLMs (Nguyen et al., 2024). Concurrently, (Wu et al., 2018) examines the relationship between error propagation and reverse direction decoding in machine translation. The second area focuses on the reversal curse (Berglund et al., 2023;

Zhu et al., 2024), where an LLM trained to understand "A is B" may struggle to generalize to "B is A". Reversing the text is proposed as a solution to this problem (Golovneva et al., 2024; Guo et al., 2024). Unlike previous works on machine translation or the reversal curse, we use reverse texts to simulate and explore reverse thinking in LLMs.

Our applications are partially related to the training data selection for LLMs, which is mainly divided into heuristic and model-based methods (Longpre et al., 2023). Heuristic methods filter out low-quality data by defining various rules, such as the ratio of nouns and verbs (Raffel et al., 2020; Penedo et al., 2023; Chowdhery et al., 2023; Sharma et al., 2024). Model-based methods filter data by training selection models or based on the perplexity of language models (Wenzek et al., 2019; Xie et al., 2023; Wettig et al., 2024). However, our data selection method is an extra bonus derived from the reverse thinking analysis.

3 Experimental Setting

Forward and Reverse Training. Given a original text, it can be represented as a sequence after tokenization, which is used for forward training. To perform reverse training, we directly reverse the original token sequence to construct a reverse training sample. While some studies explore keep the original orders of detected words or entities during reverse (Golovneva et al., 2024; Guo et al., 2024), we choose the simplest operation to avoid the various performance of detection modules in different domains and languages. The Llama2-7B (Touvron et al., 2023) (or the random initialized version) is selected as the default backbone in this paper.

Datasets. In RQ1, we used the multilingual mC4[1] (Raffel et al., 2020) dataset to compare LLMs' ability to handle forward and reverse texts under continued and from-scratch pretraining settings. In subsequent experiments, we used the carefully cleaned SlimPajama[2] (Soboleva et al., 2023a) dataset, which includes seven different source domains. Testing LLMs trained on the multilingual mC4 dataset with samples from the SlimPajama dataset can further confirm our findings are general.

[1] `https://huggingface.co/datasets/allenai/c4`

[2] We use the widely-used public sampled version for experiments: `https://huggingface.co/datasets/DKYoon/SlimPajama-6B`

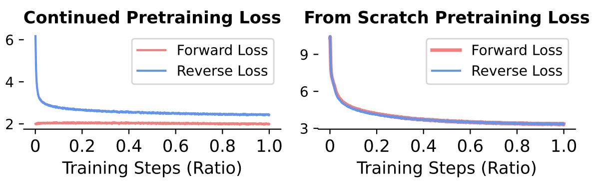

Figure 1: Pre-training loss for both continued setting and from-scratch settings in English.

4 Experiments

4.1 RQ1: *Can LLMs do reverse thinking?*

To explore LLMs' reverse thinking capabilities, we investigate two pretraining approaches: (1) continued training from a well-trained model checkpoint and (2) pretraining from scratch with random initialization. Figure 1 compares training losses (average sample losses within training batches) for English using both methods on the multilingual mC4 dataset, while Figure 5 in the Appendix shows analogous results for other languages.

In the continued pretraining setting, the forward loss for forward-thinking remains stable due to extensive training in the initial pretraining stage. In contrast, the reverse loss for reverse thinking, initially high, decreases rapidly after a few training steps. Notably, the forward loss is consistently higher than the reverse loss during continued pretraining. We speculate this occurs because the initial pretraining corpora consists entirely of forward-direction texts, imparting a natural directional bias to the LLMs. Consequently, the models find processing reverse information more challenging, similar to human difficulties with reverse thinking.

Interestingly, in the from-scratch pretraining, the loss curves for both text directions converge almost identically. This pattern, also observed in other languages, indicates that LLMs can learn to process forward and reverse-thinking inputs with similar proficiency. This is because the model learns from both forward and reverse texts simultaneously with randomly initialized parameters, avoiding the initial forward-direction bias in well-trained models.

4.2 RQ2: *Does data domain influence the ability of LLMs' reverse thinking?*

Based on the observation in Section 4.1, we focus on the from-scratch pretraining setting, where trained LLMs show almost equal losses from both forward and reverse directions when trained simultaneously. This raises the question of whether reverse loss consistently exceeds forward loss across all texts or if there are instances where reverse

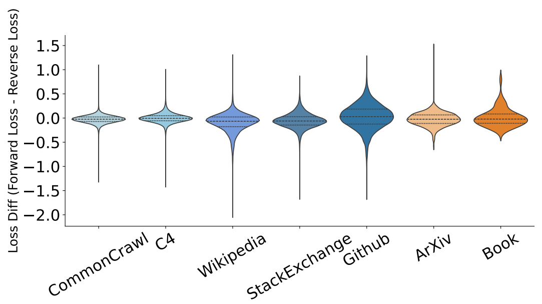

Figure 2: Loss difference distribution across domains.

learning incurs a lower loss. To explore this, we use the Slimpajama (Soboleva et al., 2023b) text dataset, which covers a broad range of domains, for case-level evaluation.

We first compute the average loss difference (Forward Loss − Reverse Loss) for each text and associate each text with its source label. The overall case-level loss difference distribution across different source domains is shown in Figure 2. Observed that the loss differences of the text samples are centered around zero, showing an approximately normal distribution. Importantly, this indicates that reverse-direction loss is not universally higher than forward-direction loss. In fact, for over half of the texts, reverse prediction of the next tokens is comparatively easier.

As depicted in Figure 2, compared to web-scraped corpora such as Wikipedia and Common Crawl, the distributions of loss differences from Book and ArXiv are generally less skewed towards easier forward-thinking. Furthermore, a larger proportion of texts in Book and ArXiv are easier to predict in the reverse direction compared to the forward direction. Considering that texts from books and academic papers are typically of higher quality than web-scraped texts, we speculate that texts, where reverse prediction is more effective, are likely to exhibit better logical coherence and fluency, indicative of their superior quality. This conjecture is also reflected in domains related to code, StackExchange, and Github. From the perspective of natural language, code often features monotonous syntax and repetitive vocabulary.

From the perspective of human forward thinking and its reflection in written texts, the forward-direction token prediction task, which involves inferring the future from the present and deducing the result from the cause, is inherently more challenging. Conversely, the reverse-direction token prediction operates from known outcomes back to

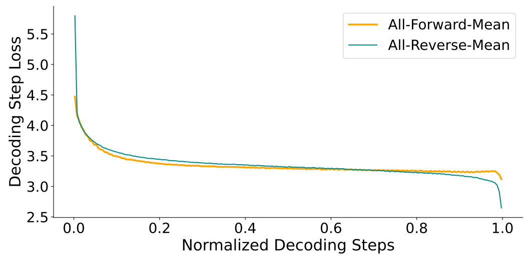

Figure 3: Mean step loss of full data during decoding.

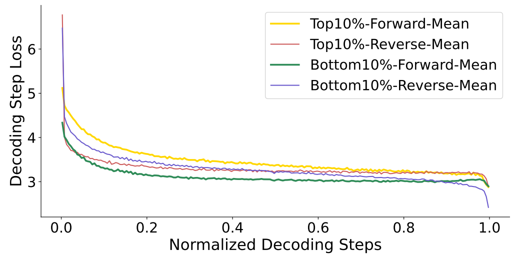

Figure 4: Mean step loss of texts with the Top-10% and Bottom-10% loss difference during decoding.

their origins, potentially simplifying the task.

4.3 RQ3: *What features make texts easier to process in the reverse direction?*

To validate our hypothesis that texts, where reverse prediction is easier than forward, correspond to higher quality, we conducted a detailed analysis of step-by-step loss changes during token decoding. We calculated and averaged the losses for each text, excluding the first and last tokens with step loss = 0 to avoid sharp changes at the start and end. To account for different input lengths, we normalized the steps of all texts to the interval $(0, 1)$. The trend of step loss changes for the entire dataset is shown in Figure 3. Overall, the reverse loss is relatively high at the initial step but decreases rapidly. In the middle step phase, the decline slope is slightly steeper for the reverse direction compared to the forward direction. In the final steps, the token prediction difficulty decreases rapidly again, while the forward loss decline trend is more stable throughout the entire decoding process.

We further examined the data with the top-10% and bottom-10% average loss differences and displayed their step loss in Figure 4. For the top-10% data, the reverse step loss quickly decreases to a level lower than the forward loss at the beginning of decoding. Conversely, for the Bottom-10% data, the reverse step loss remains higher than the forward step loss, only falling below it near the end. Table 1 summarizes the randomly selected exam-

Texts Favoring Reverse (Low Reverse Loss)	Texts Favoring Forward (Low Forward Loss)
Whether you like it or not, your garden is an open park for all of nature's creatures. ... Let's take a few minutes to learn all about ladybugs in your garden. Are Ladybugs Good for your Garden? ... Now that you know all about ladybugs and their role in controlling the aphid population, you may be interested in attracting ladybugs to your garden. ...	Ubuntu Manpage: phm2helix - calculate projections through a time varying phantom object. phm2helix - calculate projections through a time varying phantom object. phm2pj calculates projections through a time varying phantom object.

Table 1: Texts favoring reverse are often structured with clear logic flows, similar to scientific articles. Conversely, even in natural languages, texts favoring forward rely heavily on formatting to convey their sequential flow. More multilingual cases are shown in the Appendix Figure B.

Strategy	Stem	Humanities	Social Science	Other	Avg.
Original Llama2-7b	35.84	50.60	50.46	48.10	45.29
CT w/ All SlimPajama-6B	36.15	46.74	49.03	46.63	43.85
CT w/ Random 1B	35.73	46.16	48.40	47.08	43.57
CT w/ PPL Lowest Ranked 1B	36.24	45.79	47.57	45.53	43.09
CT w/ $\mathcal{S}$ Lowest Ranked 1B	34.04	45.94	45.66	42.93	41.38
CT w/ $\mathcal{S}$ Highest Ranked 1B	**37.15**	**50.93**	**50.63**	**49.82**	**46.24**

Table 2: MMLU benchmark results (Accuracy%) among different data selection strategies on LLaMA2-7b continued pre-training (CT). $\mathcal{S}$ is our proposed quality score simply computed by Forward Loss − Reverse Loss.

ples from the reverse easier and forward easier texts. The structure of the selected reverse easier cases is coherent and flows naturally, making it easy for readers to follow the information. In contrast, the forward easier texts are relatively low-quality, less coherent, and filled with repetitive phrases. This supports our earlier assumption in Section 4.2 that the reverse direction of logically coherent and well-written texts can simplify the token prediction task.

4.4 Application: *Texts favoring reverse thinking can improve original LLMs.*

As analyzed in Section 4.3, coherent and logical texts tend to have lower reverse losses compared to forward losses. Thus, given a training sample and an LLM model pretrained from scratch with both forward and reverse training, we define a simple quality score $\mathcal{S}$ using the loss difference: $\underline{\underline{\mathcal{S} = \text{Forward Loss} - \text{Reverse Loss}}}$. According to our prior analysis, A higher $\mathcal{S}$ indicates that the text, which supports reverse thinking better, signifies a high-quality sample.

To further verify this assumption, we conduct continued pretraining on the publicly released Llama2-7B. Using the SlimPajama-6B (Soboleva et al., 2023a) dataset as training data, we selected 1B samples with the lowest and highest $\mathcal{S}$ scores, respectively. The model's performance is evaluated on MMLU (Hendrycks et al., 2020) benchmark. We also compared this with the following data selection strategies: (1) **Random 1B**: randomly sample 1B data, (2) **Perplexity Lowest Ranked 1B**: select the 1B data with the lowest perplexity by the public Llama2-7B.

The results, presented in Table 2, show that the quality of training data significantly affects the performance of LLMs. Our high-quality data selection strategy ($\mathcal{S}$ Highest Ranked) outperforms other baselines, achieving the highest accuracy across various tasks on the MMLU benchmark. While using the full 6B dataset does improve over the original Llama2-7B, the improvement is marginal compared to the gain achieved through our method. This suggests that the presence of low-quality data in unfiltered training sets degrades performance, as evidenced by the significant performance decline with our low-quality selection strategy ($\mathcal{S}$ Lowest Ranked). This experiment supports the hypothesis that texts more effectively modeled by reverse thinking are of higher quality and more beneficial for LLMs in acquiring world knowledge.

5 Conclusions

In conclusion, our results demonstrate that LLMs can learn from both forward and reverse-thinking texts with comparable proficiency when trained from scratch, where human always struggle with reverse texts. This study also highlights the potential benefits of incorporating training data that favors reverse thinking. Our findings underscore the importance of exploring diverse cognitive frameworks to enhance the capabilities of LLMs.

6 Limitations

While our study demonstrates promising results in training LLMs with reverse thinking, several limitations should be acknowledged to provide a comprehensive understanding of the findings and guide future research.

Firstly, the simulation of reverse thinking by simply reversing token sequences may not fully capture the complexity and nuances of true reverse cognitive processes. This approach reduces reverse thinking to a syntactic level, potentially overlooking deeper semantic and contextual factors intrinsic to human reverse thinking.

Secondly, the evaluation metrics used in our study, such as performance on downstream benchmarks like MMLU, may not fully encompass the benefits or limitations of reverse thinking training. These metrics primarily measure specific aspects of language understanding and reasoning, potentially overlooking other critical dimensions influenced by reverse thinking, such as creativity or problem-solving skills.

Lastly, our research does not address the potential computational and resource challenges associated with training LLMs on reverse thinking texts. The increased complexity and processing demands could pose significant barriers to practical applications, particularly in resource-constrained environments.

In conclusion, while our findings offer valuable insights into the potential of reverse thinking in LLMs, addressing these limitations is crucial for advancing this line of research. Future studies should aim to develop more sophisticated methods for simulating reverse thinking, explore diverse and naturally occurring datasets, and consider a broader range of evaluation metrics to fully understand and harness the benefits of reverse cognitive processes in LLMs.

References

Benjamin Bergen and Ting Ting Chan. 2005. Writing direction influences spatial cognition. In Proceedings of the annual meeting of the cognitive science society, volume 27.

Lukas Berglund, Meg Tong, Maximilian Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. 2023. The reversal curse: Llms trained on "a is b" fail to learn "b is a". In The Twelfth International Conference on Learning Representations.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. Journal of Machine Learning Research, 24(240):1–113.

Derrick De Kerckhove and Charles J Lumsden. 2013. The alphabet and the brain: The lateralization of writing. Springer Science & Business Media.

Jian Gao and Lan Wang. 2019. Reverse thinking teaching discussion in high school information technology under new curriculum standards. In 2019 14th international conference on computer science & education (ICCSE), pages 222–226. IEEE.

Olga Golovneva, Zeyuan Allen-Zhu, Jason Weston, and Sainbayar Sukhbaatar. 2024. Reverse training to nurse the reversal curse. arXiv preprint arXiv:2403.13799.

Jiatao Gu, Qi Liu, and Kyunghyun Cho. 2019. Insertion-based decoding with automatically inferred generation order. Transactions of the Association for Computational Linguistics, 7:661–676.

Qingyan Guo, Rui Wang, Junliang Guo, Xu Tan, Jiang Bian, and Yujiu Yang. 2024. Mitigating reversal curse via semantic-aware permutation training. arXiv preprint arXiv:2403.00758.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. In International Conference on Learning Representations.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. arXiv preprint arXiv:2310.06825.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. ArXiv, abs/2001.08361.

Shayne Longpre, Gregory Yauney, Emily Reif, Katherine Lee, Adam Roberts, Barret Zoph, Denny Zhou, Jason Wei, Kevin Robinson, David Mimno, et al. 2023. A pretrainer's guide to training data: Measuring the effects of data age, domain coverage, quality, & toxicity. arXiv preprint arXiv:2305.13169.

Anh Nguyen, Nikos Karampatziakis, and Weizhu Chen. 2024. Meet in the middle: A new pre-training paradigm. Advances in Neural Information Processing Systems, 36.

Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only. arXiv preprint arXiv:2306.01116.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. Journal of machine learning research, 21(140):1–67.

Vasu Sharma, Karthik Padthe, Newsha Ardalani, Kushal Tirumala, Russell Howes, Hu Xu, Po-Yao Huang, Shang-Wen Li, Armen Aghajanyan, and Gargi Ghosh. 2024. Text quality-based pruning for efficient training of language models. arXiv preprint arXiv:2405.01582.

Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023a. SlimPajama: A 627B token cleaned and deduplicated version of RedPajama.

Daria Soboleva, Faisal Al-Khateeb, Robert Myers, Jacob R Steeves, Joel Hestness, and Nolan Dey. 2023b. Slimpajama: A 627b token cleaned and deduplicated version of redpajama.

John Sweller. 2020. Cognitive load theory and educational technology. Educational Technology Research and Development, 68(1):1–16.

Donald W Taylor and Olga W McNemar. 1955. Problem solving and thinking. Annual review of psychology, 6(1):455–482.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288.

Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán, Armand Joulin, and Edouard Grave. 2019. Ccnet: Extracting high quality monolingual datasets from web crawl data. arXiv preprint arXiv:1911.00359.

Alexander Wettig, Aatmik Gupta, Saumya Malik, and Danqi Chen. 2024. Qurating: Selecting high-quality data for training language models. arXiv preprint arXiv:2402.09739.

Lijun Wu, Xu Tan, Di He, Fei Tian, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2018. Beyond error propagation in neural machine translation: Characteristics of language also matter. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 3602–3611.

Sang Michael Xie, Shibani Santurkar, Tengyu Ma, and Percy S Liang. 2023. Data selection for language models via importance resampling. Advances in Neural Information Processing Systems, 36:34201–34227.

Long Zhou, Jiajun Zhang, and Chengqing Zong. 2019. Synchronous bidirectional neural machine translation. Transactions of the Association for Computational Linguistics, 7:91–105.

Hanlin Zhu, Baihe Huang, Shaolun Zhang, Michael Jordan, Jiantao Jiao, Yuandong Tian, and Stuart Russell. 2024. Towards a theoretical understanding of the'reversal curse'via training dynamics. arXiv preprint arXiv:2405.04669.

A Implementation

For fair comparison, we fix the learning rate as 5e-5 and the batch size as 48 for all experiments. The LLaMA2-7B (Touvron et al., 2023) is selected as the backbone. All experiments are conducted on a workstation with 8 pieces of NVIDIA A100-SXM-80GB GPUs. We conduct all experiments based on LLaMA2-7B (Touvron et al., 2023) and 8 pieces of `NVIDIA A100-SXM-80GB GPUs`. For a fair comparison, we fix the learning rate as 5e-5 and batch size as 48.

B Multilingual Experiment

We show the multilingual figure corresponding to RQ1 (Section 4.1) in Figure 5. Noted that Arabic text tokenized by LLaMA2 the same orientation as English, with tokens from the first logical sentence of a paragraph positioned on the left rather than the right.

We also randomly sample more multilingual cases showing texts favoring reverse or forward is shown in Figure B. The observation is consistent across different languages: texts that favor reverse thinking are often structured with clear logical flows, while texts that favor forward thinking rely heavily on formatting to convey their sequential progression.

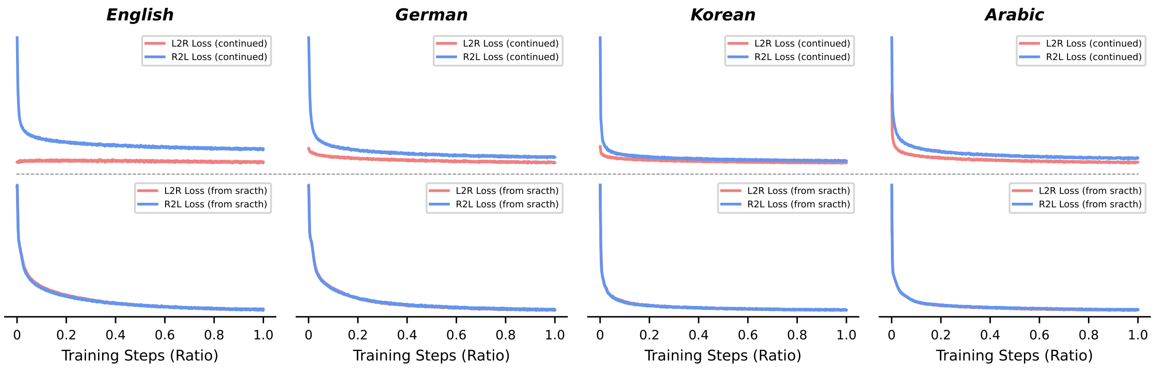

Figure 5: Pretraining loss for both continued setting and from scratch settings in four languages.

Language	Texts Favoring Reverse (Low Reverse Loss)	Texts Favoring Forward (Low Forward Loss)
German	In aller Munde, in aller Ohren – an Jonas Kaufmann kommt man derzeit nicht vorbei. Startenor, Herzensbrecher, ein echtes Münchner Kindl noch dazu, hat sich Kaufmann in die internationale erste Riege gesungen. „Seine Intensität und seine Eleganz, die Geschmeidigkeit seiner Stimme und seiner Körpersprache, kombiniert mit seiner Musikalität und seinem strahlenden Aussehen, machen ihn zum Inbegriff des Opernstars im 21. Jahrhundert“, schwärmte der Herausgeber der Opera News. Und so wird Jonas Kaufmann seit geraumer Zeit weltweit gefeiert – nicht nur an den größten Opernhäusern, sondern auch als Protagonist in Gustav Mahlers „Lied von der Erde“, als Interpret von Hugo Wolfs „Italienischem Liederbuch“ oder als leidenschaftlicher Tenor, wenn er in einer Hommage an die unsterbliche Musik Italiens ihren Evergreens eine besondere Magie verleiht. Im Rahmen einer neuen CD Produktion befindet sich das Programm für die Tournee 2019/2020 derzeit noch in der Planung und wird zu einem späteren Zeitpunkt bekannt gegeben. Die Konzertpremiere des neuen Programms wird im Oktober 2019 in Wien stattfinden. Wir dürfen also gespannt sein, welche Überraschungen uns Jonas Kaufmann nach Luzern mitbringen wird.	Die besten Yaroslavl Pauschalreisen 2018 - TripAdvisor Yaroslavl Pauschalreisen - Yaroslavl Urlaub Reisen Yaroslavl – Urlaub Yaroslavl Yaroslavl Urlaub Urlaubsangebote für Yaroslavl Spielen Sie mit dem Gedanken, eine Reise nach Yaroslavl zu buchen? Ob Sie einen Romantikurlaub, eine Familienreise oder ein All-Inclusive-Paket planen, die Pauschalreisen nach Yaroslavl auf TripAdvisor machen die Reiseplanung einfach und erschwinglich. Vergleichen Sie Hotel- und Flugpreise für Yaroslavl und finden Sie so auf TripAdvisor die perfekte Pauschalreise nach Yaroslavl. Reisende wie Sie haben 7.983 Bewertungen geschrieben und 10.284 authentische Fotos für Hotels in Yaroslavl gepostet. Buchen Sie Ihren Urlaub in Yaroslavl noch heute! Familienfreundliche Hotels in Yaroslavl “Gute Lage, ein Park und Kotorosl Ufer fußläufig gut erreichbar. Zimmer sind sauber und werden immer gut aufgeräumt. Ein sehr bequemes Bett, das man sehr selten findet. Auch einen sehr guten und ...

Language	Texts Favoring Reverse (Low Reverse Loss)	Texts Favoring Forward (Low Forward Loss)
Korean	"미국 동영상 서비스 시장, 최종 승자는 누구? - B2B IT 전문가 진행 생방송토크 웨비나 전 세계에서 인터넷 동영상 서비스(Over The Top, OTT) 경쟁이 한창이다. 글로벌 온라인 동영상 스트리밍 서비스의 선두주자 넷플릭스, 아마존닷컴의 인터넷 주문형 동영상 서비스 아마존 비디오, 동영상 공유 사이트 유튜브 등 각자의 서비스를 내세우며 피 터지는 경쟁을 하고 있다. 중심지는 아무래도 미국이다. 글로벌 IT기업의 집결지인 미국 무대를 먼저 사로잡아야 전 세계 고객들을 사로잡을 수 있다는 생각으로 오리지널 콘텐츠 개발 등 각종 공격적 마케팅 전략을 쏟아내고 있다. 콘텐츠 개발을 위한 투자 예산도 어마어마하다. 지난 4월7일 <비즈니스인사이더> 보도에 따르면, 아마존이 2017년 동영상 서비스 강화를 위해 투입할 예산이 45억달러, 우리 돈 5조1천억원 규모라는 JP모건 애널리스트들의 분석이 나왔다. 브라이언 올사브스키 아마존 CFO 역시 “아마존 비디오에 대한 투자를 두 배가까이 늘릴 것”이라고 말한 바 있다. 넷플릭스도 만만치 않다. 넷플릭스는 지난해 말, 2017년 서비스 강화를 위해 50억달러, 우리돈 5조7천억원 규모를 투입할 예정이라고 말했다. 두 회사의 투자 규모만 합쳐도 우리돈 12조원 정도 예산이니 가히 엄청나다고 할 수 있다. (자료=컴스코어) 미국 인터넷 시장조사 연구기업 컴스코어가 OTT 서비스 시장에 대한 조사 보고서를 4월10일 내놓았다. 컴스코어에 따르면, 2016년 12월을 기준으로 미국 내에 인터넷 연결망을 가진 가구 중 53%인 약 4900만 가구가 인터넷 동영상 서비스에 가입했다고 한다. 단순히 가입 규모에 그치지 않는다. 이들의 전체 평균 시청 시간은 월 평균 19일, 일 평균 2.2시간이다. 현재 미국인들의 하루 평균 TV 시청 시간은 4시간 수준이다. 케이블 위성 방송으로만 TV를 시청하던 전통적인 시청 패턴이 완전히 변화하고 있음을 알 수 있다. ...	"[37% 세일] Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 쿠폰 코드 인기 쿠폰, Jul 2020 - iVoicesoft 인기 쿠폰 › Cdkeys 쿠폰 코드 2020 › Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 쿠폰의 할인 할인 코드 37% 세일, 여름 제공 간단히 버튼을 클릭하십시오 [할인 된 가격으로 즉시 구매] 쿠폰을 사용하려면 37% 할인 코드. 쿠폰 코드가 포함되었습니다. 결제시 코드를 입력하십시오. 특별 승진의 (16.42Φ″)16.42 절약 여름 은 위대하 Cdkeys 제공 받기에 완벽한시기입니다. 2020 년 여름 제공 위해 지금 청구하십시오. 현재 거래: 37% 할인 Star Wars Battlefront II 2 - Celebration Edition Xbox One (US). Cdkeys 에서 원하는 것을 가져올 수있는 최고의 기회. 제한된 시간 동안 만. 결제시 코드를 입력하십시오. Cdkeys 쿠폰 코드: 최고의 세일즈 프로모션 사용하여 매력적인 가격으로 훌륭한 제품을 찾으십시오. 37% 할인 Star Wars Battlefront II 2 - Celebration Edition Xbox One (US), 16.42 절약. 쇼핑하려면 클릭하세요. 제한된 시간 동안 만. Star Wars Battlefront II 2 - Celebration Edition Xbox One (US)에 대하여 Star Wars Battlefront II 2 - Celebration Edition Xbox One (US)소개 Get 37% OFF of Star Wars Battlefront II 2 - Celebration Edition Xbox One (US), a 위대하 in 여름 제공 Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 쿠폰 코드. Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 위대하 여름 제공 37% 쿠폰 코드. 왜 우리의 Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 쿠폰 코드를 적용해야합니까? 간단 해! 최신 Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 프로모션 코드를 수집하여 제공했습니다, 가장 큰 할인으로. 또한 모든 Cdkeys 제품에 대한 최고의 절감 효과를 제공합니다. Star Wars Battlefront II 2 - Celebration Edition Xbox One (US) 할인 코드에 대한 의견"