

Acquiring Linguistic Knowledge from Multimodal Input

Theodor Amariucaí
ETH Zürich
tamariucaí@ethz.ch

Alex Warstadt
ETH Zürich
awarstadt@ethz.ch

Abstract

In contrast to children, language models (LMs) exhibit considerably inferior data efficiency when acquiring language. In this submission to the BabyLM Challenge (Warstadt et al., 2023), we test the hypothesis that this data efficiency gap is partly caused by a lack of multimodal input and grounding in the learning environment of typical language models. Although previous work looking into this question found that multimodal training can even harm language-only performance, we speculate that these findings can be attributed to catastrophic forgetting of complex language due to fine-tuning on captions data. To test our hypothesis, we perform an ablation study on FLAVA (Singh et al., 2022), a multimodal vision-and-language model, independently varying the volume of text and vision input to quantify how much text data (if any) can be offset by vision at different data scales. We aim to limit catastrophic forgetting through a multitask pretraining regime that includes unimodal text-only tasks and data sampled from WiT, the relatively diverse Wikipedia-based dataset (Srinivasan et al., 2021). Our results are largely negative: Multimodal pretraining does not harm our models’ language performance but does not consistently help either. That said, our conclusions are limited by our having been able to conduct only a small number of runs. While we must leave open the possibility that multimodal input explains some of the gap in data efficiency between LMs and humans, positive evidence for this hypothesis will require better architectures and techniques for multimodal training.

1 Introduction

Children can learn language from a relatively small amount of linguistic input: at most 100 million words (Gilkerson et al., 2017). By contrast, the quantity of training data a language model needs to achieve strong grammar and language performance is on the order of billions or tens of billions

of words (Zhang et al., 2021). This data efficiency gap may be due, in part, to innate differences in learning mechanisms between models and humans, but environmental differences likely play a role as well (Warstadt and Bowman, 2022). This work tests the hypothesis that the lack of visual grounding in language models accounts for some of the gap in data efficiency.

The likelihood of finding evidence for this hypothesis rests largely on two factors: (1) its cognitive plausibility and (2) its technological viability. If vision does help children learn language, then there ought to be some way of incorporating vision into text-only language models that improves their learning ability. However, whether or not we can find this approach depends on the present technological capabilities of multimodal models.

To address the first point, one cognitively-motivated mechanism for how children integrate nonlinguistic sensory data in language learning is cross-situational learning (XSL) (Smith and Smith, 2012). This theoretical mechanism holds that the learner accumulates statistical evidence about word meanings by observing multiple instances of co-occurring word-object pairs across many different real-world situations (Smith et al., 2011; Kachergis et al., 2014; Zhang et al., 2019). Encouragingly, Nikolaus and Fourtassi (2021) find that, in a highly constrained visual-linguistic domain, computational multimodal models *do* benefit from cross-situational learning.

To address the second point, prior evidence that vision will improve language models given current technologies is, at best, mixed. Recent approaches have successfully trained Transformer-based multimodal language models using self-supervised objectives resembling those developed originally for the training of unimodal models (Tan and Bansal, 2019). Nevertheless, in comparison to the unimodal models, multimodal LMs often perform relatively poorly on language-only tasks (Iki and

Aizawa, 2021). We hypothesize that these shortcomings may be due to the common practice of training multimodal models by fine-tuning pre-trained language models on captions data.

Our approach addresses these limitations in two ways that differ from most previous pretraining recipes. First, we use the FLAVA architecture and follow its multitask training procedure (Singh et al., 2022). Second, we train on the Wikipedia-based WiT dataset (Srinivasan et al., 2021), which pairs images with a mixture of strongly aligned (but formulaic) captions and weakly aligned (but syntactically complex) articles. Despite these efforts, our results show that the addition of image data and multimodal training objectives leads to no reliable improvement over text-only baselines on benchmarks for grammar (BLiMP; Warstadt et al., 2020a), understanding (GLUE; Wang et al., 2018), and generalization (MSGs; Warstadt et al., 2020b). We conclude that, to the extent that multimodality is partly responsible for the data-efficiency gap, present multimodal (and multitask) pretraining methods do not benefit from this richer learning signal.

To summarize, this work brings forward three main contributions:

1. We develop a robust codebase¹ for pretraining (from scratch) large multimodal LMs under varying text and vision input configurations.
2. We evaluate, in this controlled environment, the effects of the visual signal on the model’s textual encoder (hence, its linguistic ability).
3. We investigate plausible mechanisms for how incorporating visual input into the pretraining procedure might affect linguistic behavior.

2 Background

Prior work on multimodal language model training can be roughly differentiated by whether the main objectives are cognitively-oriented or engineering-oriented. So far, neither of these directions has produced clear evidence supporting the hypothesis that multimodality aids language learning at the scale of human language acquisition. Many cognitively oriented contributions are limited by a small data scale or a restricted domain. By contrast, engineering-oriented contributions using state-of-the-art Transformer-based architectures achieve

¹<https://github.com/amariucaitheodor/acquiring-linguistic-knowledge>

more developmentally plausible scale and diversity but emphasize multimodal performance over language learning.

2.1 Cognitively Oriented Approaches

Infants enter a diverse and abundant visual world where they develop mental models to comprehend and mimic the patterns they encounter. These mental models empower them to grasp and anticipate their surroundings and accomplish objectives by incrementally improving their communicative abilities (Roy and Pentland, 2002). Relatedly, the impact of vision on specific aspects of human language learning has been an important question in human development for decades.

Contemporary research has tried to answer this question through computational simulations of cognitive processes involved in language acquisition. Multimodal models trained on visual question answering or reference games can use cross-situational learning to learn grounded meanings of words (Mao et al., 2019; Wang et al., 2021; Nikolaus and Fourtassi, 2021; Portelance et al., 2023). Nonetheless, computational models show different learning biases than humans in many cases, at least in the absence of specific training or architectural interventions (Gauthier et al., 2018; Vong and Lake, 2022). Ultimately, however, all of these studies are limited in their cognitive plausibility and language learning by a reliance on supervised training on small, artificial datasets in which texts and images correspond to arrangements of a limited set of objects in a simple, usually static scene.

Other studies aim for more naturalistic training. Lazaridou et al. (2017) and Chrupała et al. (2015) are notable for pioneering self-supervised training objectives for multimodal models several years before the advent of Transformer architectures trained on masking objectives. Wang et al. (2023) train LMs on data from the SAYCam dataset (Sullivan et al., 2021), pairing (written) child-directed utterances with visual data from the child’s point of view. While this data domain is nearly ideal from a developmental plausibility perspective, the available data is too small to model anything past the first month of development.

Finally, we note that most of the studies in this area focus primarily on word learning. However, the data efficiency gap applies more broadly to language learning. Recent studies evaluating contemporary Transformer-based models have largely re-

ported negative results for the effect of multimodality on semantics (Shahmohammadi et al., 2022), commonsense reasoning (Yun et al., 2021), and learning biases (Kuribayashi, 2023). To the best of our knowledge, ours is the first work to perform targeted syntactic evaluation (Marvin and Linzen, 2018; Warstadt et al., 2020a; Hu et al., 2020) on multimodal models.

2.2 Engineering Oriented Approaches

The most effective recent approaches for training multimodal language models generalize the self-supervised objectives (and Transformer-based architectures) that have become dominant for unimodal language models such as BERT (Devlin et al., 2019) to the vision-and-language domain (Li et al., 2020, 2019; Zhou et al., 2020; Tan and Bansal, 2019; Chen et al., 2020; Yu et al., 2021; Lu et al., 2019; Bugliarello et al., 2021; Pramanick et al., 2022).

These studies, for the most part, share many aspects of a typical recipe: First, they initialize all or some of the model parameters with the pretrained weights of a model such as BERT. Second, they fine-tune (using one or more self-supervised objectives) on a dataset of image-caption pairs.² Finally, the model is evaluated on multimodal tasks such as visual question answering or image captioning.

While the ability to perform such grounded tasks is the key advantage of multimodal models over unimodal ones, it is critical for our research question to examine whether this advantage comes at the cost of language ability. Unfortunately, few of the works that train new multimodal models evaluate on language-only tasks. Some works perform this evaluation post hoc. Iki and Aizawa (2021) study five multimodal architectures, all initialized with BERT and fine-tuned using identical data and training objectives by Bugliarello et al. (2021). Evaluating on the GLUE benchmark (Wang et al., 2018), they find that, in nearly all cases, the original pretrained BERT outperforms the models with additional multimodal fine-tuning. Similar results are reported by Madasu and Lal (2023) and Yun et al. (2021).

²For example, in the masked multimodal modeling task (MMM; Tan and Bansal, 2019), regions of an aligned image-text pair are randomly masked before being input into the model and then predicted. As the information from the image presumably helps reconstruct the masked text (Frank et al., 2021), this objective encourages learning text representations that encode information from the visual modality (and vice-versa).

From a human development perspective, it may seem unintuitive that additional supervision on images harms language performance. However, from a machine learning perspective, this finding is easy to explain as an example of domain mismatch (Yun et al., 2021), catastrophic forgetting (McCloskey and Cohen, 1989), under-parameterization (Amaricai, 2023), or other similar technical reasons.

BERT’s original training data (Wikipedia and books) is diverse in terms of writing style and subject matter. By contrast, captions datasets commonly used to train multimodal LMs, such as MS COCO (Chen et al., 2015) or Visual Genome (Krishna et al., 2017), consist entirely of short formulaic physical descriptions of objects or scenes. Hence, the text domain that the models were trained on most recently bears little resemblance to the texts in the GLUE tasks, for example. Furthermore, the multimodal tasks incentivize using the models’ limited parameters for both text *and* image processing, potentially sacrificing language ability.

Our experiments, which we describe in Section 3, are designed to address these issues through two complementary approaches: First, we prevent catastrophic forgetting by multitask-training on the language-only masked language modeling (MLM) objective jointly with the multimodal objectives. Second, we lessen the impact of domain mismatch by training on data that pairs images, not just with captions, but also with longer and more complex texts.

3 Methods

We conduct experiments to uncover differences in how language models’ linguistic abilities change as the amount of visual input varies. We pretrain and evaluate multimodal LMs in eight conditions, derived by independently varying the volume of text data (10M or 100M words) and image data (none, 40K, 400K, or 4M images). We perform only one training run for each of the eight conditions due to computing constraints (see Limitations, Section 5). The text quantities are compatible with both human-scale linguistic exposure (Gilkerson et al., 2017) and the BabyLM strict-small and strict tracks (Warstadt et al., 2023).

3.1 Dataset

All the data for our experiments comes from WiT, a large, multimodal dataset entirely sourced from Wikipedia (Srinivasan et al., 2021). Our choice

of WiT was motivated by its size and the diversity and complexity of its text data. English WiT includes 5.5M image–text pairs,³ making it one of the largest public datasets of its kind. It contains extended passages from Wikipedia articles, offering a more representative sample of sentence types than typical multimodal datasets sourced from captions. Furthermore, WiT features multiple types of text aligned with a given image. From most strongly aligned to most weakly aligned, these include alt text, captions, article text from the same section as the image, and article text from the lead section. Together with the fact that Wikipedia covers many different concepts and real-world entities, we hypothesize that WiT provides an adequately rich environment for supporting cross-situational learning while maintaining strong grammar and language understanding performance.

We subsample from the English portion of WiT to reach the desired data volume for each modality. For training purposes, we use either one (when either modality is 0%) or three (when both modalities are non-zero) data loaders. For example, when training on 100M words and 40K images, we sample the first 10% of the pairs for the text unimodal data loader, the first 1% for the vision unimodal data loader, and the first 1% = $\min(10\%, 1\%)$ for the multimodal data loader (containing paired images and text). Hence, all images in this configuration will be paired with some text, but not all texts will be paired with an image. This logic also implies that some images and texts will be seen both in the multimodal and their corresponding unimodal data loaders.

3.2 Model

For our experiments, we use the FLAVA model architecture and training objectives (Singh et al., 2022). We choose to study FLAVA for two reasons: First, Singh et al. (2022) conduct a controlled comparison between a unimodally trained FLAVA text encoder and a fully multimodal FLAVA, and they report improved performance on language-only tasks from the multimodal model. As such, FLAVA is the only example of a large multimodal model for which prior (anecdotal) evidence supports our hypothesis that vision can help language learning. Second, FLAVA is trained in a multi-task setting on a combination of unimodal text,

³WiT is also multilingual, containing over 30M pairs in over 100 languages. During preprocessing, however, we only sample English text.

unimodal vision, and multimodal objectives. This methodology addresses our concern (Section 2.2) that other common multimodal training recipes can lead to catastrophic forgetting of linguistic ability.

FLAVA’s architecture combines three modality-specific encoders: Text and vision embeddings are fed into unimodal text and vision encoders, respectively, and the hidden states output by these encoders are concatenated before being fed into a multimodal encoder. For the unimodal objectives, task-specific heads can be placed after the corresponding unimodal encoder. All encoders are based on the ViT-B/16 encoder (Dosovitskiy et al., 2021).⁴

Following the original work, we pretrain models from scratch using multitask learning with the following five objectives: masked image modeling, masked language modeling, masked multimodal modeling for both text and vision, image-text matching, and cross-modal contrastive learning. More details on each objective, as well as the encoder architecture itself, can be found in the original paper (Singh et al., 2022).

3.3 Training Details

Hyperparameters We perform a hyperparameter search and empirically settle on the following values:

- Warmup steps: 10^4
- Batch size: 4096 effective = 32×2 GPUs \times 64 accumulation steps
- Learning rate (text encoder): 7.5×10^{-4}
- Learning rate (other encoders): 10^{-3}
- Precision: bf16 mixed
- Seed: 5501650
- Adam optimizer:
 - Epsilon: 10^{-8}
 - Weight decay: 0.1
 - Betas: [0.9, 0.999]

We use two distinct learning rates because a lower value is commonly recommended for text-only pretraining (Liu et al., 2019; Devlin et al., 2019), while a higher one was originally used for

⁴*B/16* refers to a base-sized architecture with 86M total parameters using a patch resolution of 16x16. We opt for this version of the encoder based on the authors’ observation that ViT-B/16 performs just as well as the larger alternatives when pretraining on smaller datasets of under 300M images (such as WiT).

multimodally pretraining FLAVA. Multiple strategies for correctly choosing modality-specific learning rates are treated extensively in Yao and Mihalcea (2022), where the "Keep" Strategy (ours) is among the most straightforward of them. While simple, it outperforms (in the authors' empirical study) the global learning rate strategy as it ensures that each unimodal subpart still has effective gradients when training the fusion model.

Software We use Pytorch Lightning (Falcon and The PyTorch Lightning team, 2019) as the main training framework and Weights and Biases (Biewald, 2020) to track relevant metrics in real time. We use the Huggingface datasets library (Lhoest et al., 2021) to interleave the modality-specific datasets, and the HuggingFace transformers library (Wolf et al., 2020) to access and train randomly-initialized FLAVA models.

Hardware We run each training job in Distributed Data-Parallel mode, across two NVIDIA Tesla A100 Ampere 40 GB graphics cards, on the same node, in ETH Zürich's Euler datacenter. For each of the two GPUs, there are 4 CPU workers loading data (this number was empirically found to be optimal), with each CPU worker having 10GB of RAM available. The average runtime for our jobs running on 100M words was six days, and for 10M words, it was three days. Thus, we count a total of (2 GPUs) * (8 jobs) * (108 hours / job) = 1728 GPU hours used to train the models reported in this study, not counting our hyperparameter search.

Dataloader Sampling Weights During multimodal pretraining, we alternate samples from three data loaders with independent weights, initialized (and normalized) proportional to their sizes. For example, for the condition with 100M words and 40K images (hence, all images and 10M words of text are paired), we would have the following initial sampling weights: 0.833 (text), 0.083 (vision), 0.083 (multimodal). For maximal text encoder performance, we perform a hyperparameter search and determine a simple rule-based approach to further improve the distribution of the sampling weights: If text is not the predominant modality, we change it to the uniform distribution; otherwise, we leave the initial weights unchanged.

Modality-Specific Early Stopping We develop custom logic to prevent the models from overfitting on any given modality. For example, when the

Words \ Images	Images			
	None	40K	400K	4M
10M	4604	5346	4876	4876
100M	23818	12267	16344	16542

Table 1: Model checkpoints (training step #) chosen for evaluation based on the masked language modeling validation loss (Figure 4).




Figure 1: PPPL performance for the two data volumes of 10M and 100M words. The training steps on the x-axis are counted across all objectives.




Figure 2: BLiMP performance for the two data volumes of 10M and 100M words. The training steps on the x-axis are counted across all objectives.

Tasks	10M Words				100M Words			
	None	40K	400K	40M	None	40K	400K	40M
(Super)GLUE (Acc., F1, MCC)	65.49	64.68	65.17	<u>65.76</u>	<u>70.42</u>	69.24	68.9	69.07
BLiMP (Acc.)	63.96	63.98	63.31	<u>64.53</u>	71.32	70.45	<u>71.9</u>	70.93
MSGS (MCC)	-12.88	-12.16	<u>-8.84</u>	-18.62	-8.66	<u>-6.18</u>	-7.41	-7.47

Table 2: Performance for each of the eight models on the BabyLM test suites (detailed version in Table 3).

sampling rate is 0.083 for the multimodal and text data loaders yet 0.833 for the vision data loader, the former two modalities will likely begin to overfit well before the latter. To avoid this, we detect increases in validation loss and (each time) halve the corresponding task’s sampling weight. If the validation loss continues to steadily increase after three validation steps, we set the task weight to 0. To prevent catastrophic forgetting of the multimodal input, we allow the models to restart training on vision and multimodal data after a certain period of inactivity (here, 10 validation phases).

Model Selection For each of the eight input configurations, we select the best model checkpoints for evaluation based on the lowest recorded masked language modeling loss on the validation set (see Figure 4 for validation losses for every training objective and model). Table 1 shows the number of training steps for the selected checkpoints from each configuration. For additional information, we also regularly evaluate the models’ pseudo-perplexity on a held-out test set (see Figure 1).

4 Results

We evaluate the selected checkpoints from all eight training configurations on the BabyLM evaluation pipeline (Warstadt et al., 2023), including evaluations on benchmarks for grammar (BLiMP; Warstadt et al., 2020a), language understanding (GLUE and SuperGLUE; Wang et al., 2018, 2019), and linguistic generalization (MSGS; Warstadt et al., 2020b). For BLiMP and pseudo-perplexity (Wang et al., 2018), we also report intermediate results for all of the training checkpoints.

Our overall results in Table 2 largely confirm earlier work finding that vision is, at best, not consistently helpful to language performance. With a data volume of 10M words, FLAVA does sometimes perform marginally better on grammar-oriented tasks in the presence of visual cues. For other evaluations and with a data volume of 100M words, we also find no consistent advantages in our experimental

setting. Of those improvements we do observe, our tests deem it unlikely that they are due to cross-situational learning (see Section 4.4).

4.1 Pseudo-perplexity

For validation, Figure 1 shows the pseudo-perplexity (PPPL; Wang and Cho, 2019) per token on a held-out evaluation subset of WIT throughout training. Unsurprisingly, PPPL is lower (better) for the 100M word models compared to the 10M word models. Additionally, the metric appears to converge for the 10M word models, while it may still be decreasing for 100M word models.⁵ The most unexpected finding is that PPPL is consistently worse as the amount of image data increases for a given amount of text data throughout training. This degradation may suggest that our multitask training procedure causes the models to sacrifice MLM performance in favor of other objectives as the proportion of visual and multimodal samples increases.

4.2 Grammaticality

We evaluate linguistic knowledge using BLiMP (Warstadt et al., 2020a), which tests the ability of models to distinguish grammatical sequences from minimally different ungrammatical ones in a zero-shot setting.

Table 2 shows the overall BLiMP performance from each condition. We notice that text quantity makes a big difference in performance. Changes in vision, on the other hand, are associated with small amounts of variation that are sometimes positive or negative. Hence, due to the lack of a consistent pattern and the small number of runs, we cannot confidently conclude that vision causes an increase or decrease in performance.

Figure 2 shows the BLiMP results for each val-

⁵Our scheduler triggered early stopping based on validation loss despite the apparent possibility that longer training might have been beneficial. More generally, there are many potential improvements to task scheduling and early stopping for multitask learning that we leave to future work.

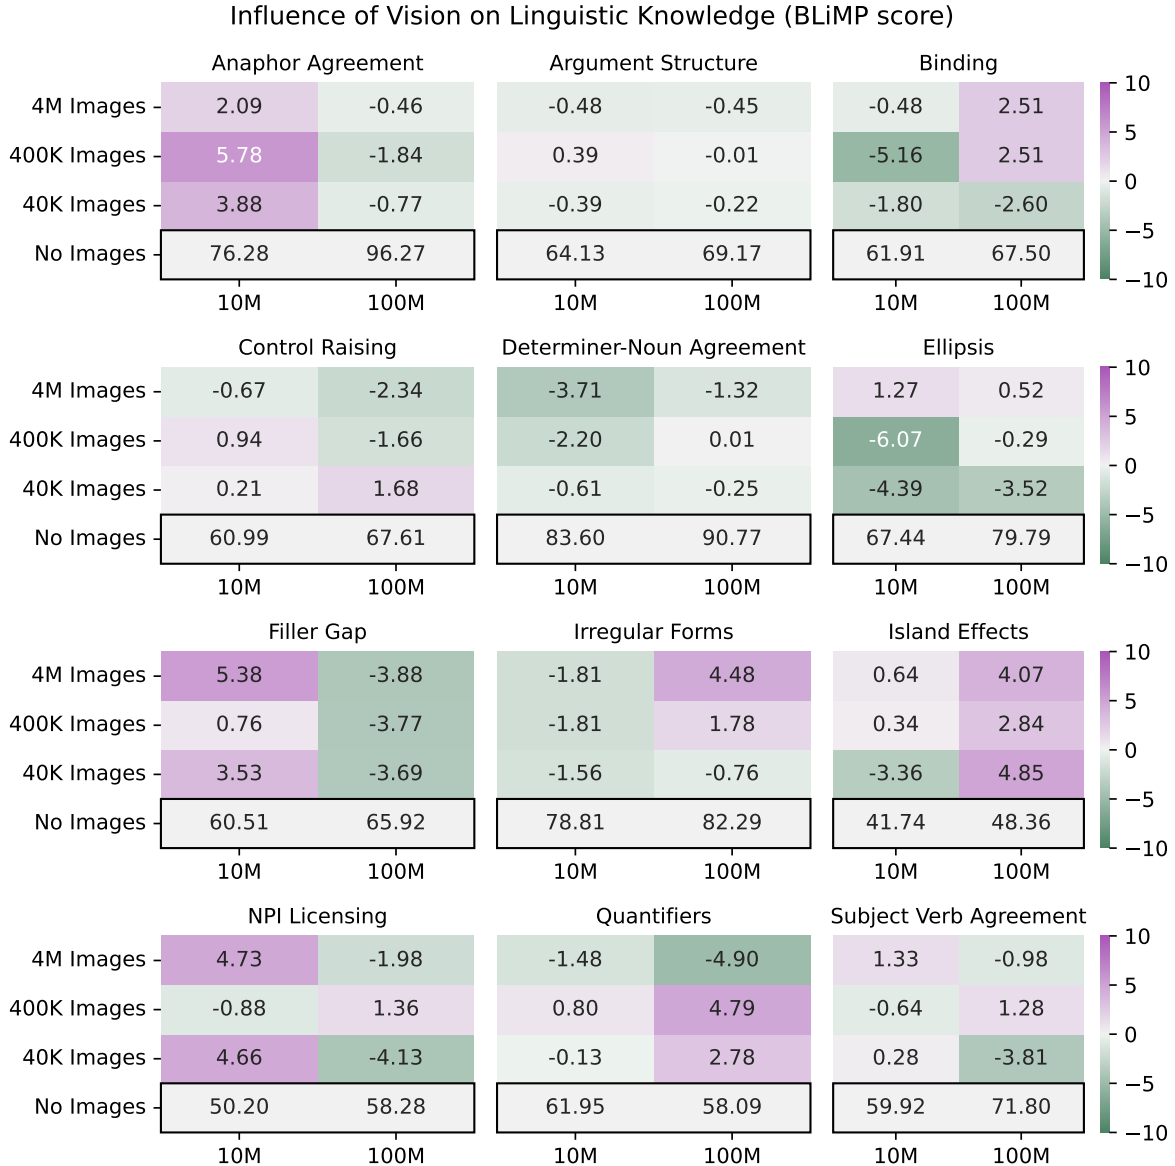


Figure 3: Zero-shot accuracies, in percentages, obtained on the BLiMP task for each grammatical category (x12) and FLAVA run configuration of input text volume (10M and 100M words) and input vision volume (0, 40K, 400K and 4M images). The model checkpoints used to generate these results were selected as described in Table 1.

idation step throughout training. For most of the duration of training, particularly for the 100M word models, models with less image data perform better. This behavior mostly matches the pattern we observe for pseudo-perplexity, except that the differences seemingly disappear by the end of training. This result confirms earlier findings that (pseudo-)perplexity is not entirely predictive of grammatical knowledge (Hu et al., 2020).

Individual BLiMP categories are more closely compared in Figure 3. Previous work shows that phenomena related to agreement have the steepest learning curves at the 10M word scale (Zhang et al.,

2021). Therefore, if the hypothesis that vision accelerates LM learning is correct, we might expect to see the greatest signs of improvement for 10M word models on this subset of test suites. Figure 3, however, shows conflicting and inconclusive results, with improvements in *anaphor agreement* but a slight degradation for *determiner-noun agreement*, and little change for *subject-verb agreement*.

We observe that multimodal pretraining may have a regularizing effect at smaller data scales: BLiMP performance improves at times although the pseudo-perplexity (i.e., test loss) is consistently higher (by 1-3 units) for the vision-infused mod-

els. Moreover, the vision-infused models run for almost twice as long before starting to overfit (8k v.s. 4k steps), gaining accuracy in areas such as *anaphor agreement*, *filler gap dependencies*, and *NPI licensing*, although not so on test suites such as *argument structure* and *subject-verb agreement*.

4.3 Fine-Tuning Evaluations

In addition to the above, we also use GLUE/SuperGLUE (Wang et al., 2018, 2019) and MSGS (Warstadt et al., 2020b) to fine-tune and evaluate all eight models on a selection of downstream tasks that focus on language understanding and linguistic generalization.

As expected, the results in Table 2 show that overall GLUE performance increases (by around 5%) at higher text data scales. Within each of the two text volume groups, however, there is no reliable improvement due to the addition of vision, though vision-infused models appear to be slightly better (relatively) at lower data scales, of up to 10M words. Generally, the models perform similarly on the selected downstream tasks (performance after fine-tuning), in line with BLiMP results.

Scores on MSGS are negative for all models, for all ambiguous subtasks (i.e., those subtasks not in the control condition), as shown in Table 3. This indicates that all of our models are consistently biased towards generalizing based on shallow surface cues rather than linguistic features.

4.4 Cross-Situational Learning

To assess the *symbolic grounding* of our models, for every input configuration checkpoint in Table 1, we evaluate the multimodal text retrieval zero-shot accuracy on ImageNet-1k (Russakovsky et al., 2015). The goal is to select, for every given *query* image, the best-fitting text caption from a pool of 1000 options. To this end, we compute cosine similarities as matching scores between the queried image’s representation and the representations of 1000 template-averaged⁶ potential captions. Lastly, we retrieve the text caption with the highest matching score for each image query. We follow Radford et al. (2021) to calculate the zero-shot accuracy.

As a baseline, we assess FLAVA pretrained on the PMD corpus and obtain *top1* and *top5* accuracies of 32% and 60%, respectively. The models we pretrain, however, obtain average *top1* and

⁶Since the captions for even a specific entity can vary, e.g., *a doodle of a car*, *a photo of a large car*, etc. we compute an average over ~ 80 such templates for each entity.

top5 accuracies of 0.1% and 0.5%, respectively. Some of the possible factors responsible for this random guess performance could be: 1) the multi-task scheduler described in Appendix 3.3 was misconfigured (this aligns with findings in Amariucaí (2023)), 2) the smaller magnitude of the training data (WiT is a subset of PMD), 3) the weak alignment between some of the text (full paragraphs) and the corresponding images (Imagenet-1k only evaluates caption alignments), or 4) the fact that we do not pretrain the vision encoder unimodally on ImageNet-1k, as in Singh et al. (2022).

5 Conclusion

We perform an ablation study on a state-of-the-art, multimodal language model under varying text and vision configurations. Our training recipe avoids the problem of catastrophic forgetting of complex language, which previous approaches fell prey to, by performing multitask training on both multimodal and unimodal tasks in a more diverse domain. Nonetheless, our results largely confirm earlier work finding that vision is (at best) not consistently helpful to language performance. During pretraining at the small 10M word scale, the FLAVA architecture (Singh et al., 2022) does sometimes appear to perform marginally better on grammar-oriented tasks in the presence of visual cues. However, for other evaluations and with a data volume of 100M words, we find no consistent advantages in our experimental setting.

At the small data scales that we pretrain our models in this study (up to 100M words and the corresponding images), our tests in Section 4.4 deem it unlikely that the models are benefiting from cross-situational learning. Alternatively, the extra parameters in the multimodal encoder could simply be increasing FLAVA’s modeling capacity, a hypothesis that we leave for future work. Regardless, multimodal pretraining seems to exhibit a regularizing effect: although pseudo-perplexity is consistently worse for the vision-infused models, grammatical performance fluctuates and is often at least as good.

We conclude that the lack of visual input alone does little to explain the large data efficiency gap between LMs and humans observed in grammar learning, though we leave open the possibility that this conclusion will change with better architectures and techniques for integrating vision and language at training time.

Limitations

The robustness of the observations made in this report is limited by the fact that each configuration (text/vision input volume) was only run once. Future work should provide at least 5 re-runs per configuration (with different seeds), as there can be considerable variance even in different models with the same configuration (McCoy et al., 2020). Due to the computational intensity of performing re-runs, this was not possible in time for this submission.

Significant GPU resources are required to effectively train large language models, partly because of the large batch sizes and the scale of the datasets. In this work, we use ≈ 1728 GPU hours on very recent hardware (further details in Section 3.3).

Finally, there is an architectural difference between the unimodal and multimodal models in our experiments. The unimodal models are trained entirely without the visual or multimodal encoders. Although these parameters are not used by the multimodal model during evaluation on language-only tasks, they are used during training, and so they may have an indirect effect on what the language encoder learns. To test whether the potential performance improvements in grammaticality and language understanding can indeed be attributed to the visual cues or rather simply to the increased number of parameters in the multimodal encoder, future work should conduct additional baseline experiments, e.g., where the images are replaced with random noise pixels.

References

- Theodor Amariuca. 2023. [Acquiring linguistic knowledge from multimodal input](#). Master thesis, ETH Zürich, Zürich.
- Lukas Biewald. 2020. [Experiment tracking with weights and biases](#). Software available from wandb.com.
- Emanuele Bugliarello, Ryan Cotterell, Naoaki Okazaki, and Desmond Elliott. 2021. [Multimodal pretraining unmasked: A meta-analysis and a unified framework of vision-and-language BERTs](#). *Transactions of the Association for Computational Linguistics*, 9:978–994.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. 2015. [Microsoft COCO Captions: Data Collection and Evaluation Server](#). In *European conference on computer vision*, pages 740–755. Springer. ArXiv: 1504.00325.
- Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. [UNITER: UNiversal Image-Text Representation Learning](#). In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX*, page 104–120, Berlin, Heidelberg. Springer-Verlag.
- Grzegorz Chrupała, Ákos Kádár, and Afra Alishahi. 2015. [Learning language through pictures](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 112–118, Beijing, China. Association for Computational Linguistics.
- 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, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *International Conference on Learning Representations*.
- William Falcon and The PyTorch Lightning team. 2019. [PyTorch Lightning](#).
- Stella Frank, Emanuele Bugliarello, and Desmond Elliott. 2021. [Vision-and-language or vision-for-language? on cross-modal influence in multimodal transformers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9847–9857, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jon Gauthier, Roger Levy, and Joshua B. Tenenbaum. 2018. [Word learning and the acquisition of syntactic-semantic overhypotheses](#). In *Proceedings of the 40th annual meeting of the cognitive science society*, Madison, Wisconsin.
- Jill Gilkerson, Jeffrey A. Richards, Steven F. Warren, Judith K. Montgomery, Charles R. Greenwood, D. Kimbrough Oller, John H. L. Hansen, and Terrance D. Paul. 2017. [Mapping the Early Language Environment Using All-Day Recordings and Automated Analysis](#). *American Journal of Speech-Language Pathology*, 26(2):248–265.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. [A Systematic Assessment of](#)

- Syntactic Generalization in Neural Language Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Taichi Iki and Akiko Aizawa. 2021. Effect of Visual Extensions on Natural Language Understanding in Vision-and-Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2189–2196, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- George Kachergis, Chen Yu, and Richard M. Shiffrin. 2014. Cross-situational word learning is both implicit and strategic. *Frontiers in Psychology*, 5.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. 2017. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. *International Journal of Computer Vision*, 123(1):32–73.
- Tatsuki Kuribayashi. 2023. Does Vision Accelerate Hierarchical Generalization of Neural Language Learners? ArXiv:2302.00667 [cs].
- Angeliki Lazaridou, Marco Marelli, and Marco Baroni. 2017. Multimodal Word Meaning Induction From Minimal Exposure to Natural Text. *Cognitive Science*, 41(S4):677–705.
- Quentin Lhoest, Albert Villanova del Moral, Patrick von Platen, Thomas Wolf, Mario Šaško, Yacine Jernite, Abhishek Thakur, Lewis Tunstall, Suraj Patil, Mariama Drame, Julien Chaumond, Julien Plu, Joe Davison, Simon Brandeis, Victor Sanh, Teven Le Scao, Kevin Canwen Xu, Nicolas Patry, Steven Liu, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Nathan Raw, Sylvain Lesage, Anton Lozhkov, Matthew Carrigan, Théo Matussière, Leandro von Werra, Lysandre Debut, Stas Bekman, and Clément Delangue. 2021. Datasets: A Community Library for Natural Language Processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184. Association for Computational Linguistics.
- Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. 2020. Unicoder-VL: A universal encoder for vision and language by cross-modal pre-training. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:11336–11344.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A simple and performant baseline for vision and language.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Avinash Madasu and Vasudev Lal. 2023. Is multimodal vision supervision beneficial to language?
- Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua Tenenbaum, and Jiajun Wu. 2019. The Neuro-Symbolic Concept Learner: Interpreting Scenes, Words, and Sentences from Natural Supervision. In *Proceedings of ICLR*.
- Rebecca Marvin and Tal Linzen. 2018. Targeted Syntactic Evaluation of Language Models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202.
- Michael McCloskey and Neal J. Cohen. 1989. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. In *Psychology of Learning and Motivation*, volume 24, pages 109–165. Elsevier.
- R. Thomas McCoy, Junghyun Min, and Tal Linzen. 2020. BERTs of a feather do not generalize together: Large variability in generalization across models with similar test set performance. In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 217–227, Online. Association for Computational Linguistics.
- Mitja Nikolaus and Abdellah Fourtassi. 2021. Evaluating the acquisition of semantic knowledge from cross-situational learning in artificial neural networks. In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 200–210, Online. Association for Computational Linguistics.
- Eva Portelance, Michael C. Frank, and Dan Jurafsky. 2023. Learning the meanings of function words from grounded language using a visual question answering model. ArXiv:2308.08628 [cs].
- Shraman Pramanick, Li Jing, Sayan Nag, Jiachen Zhu, Hardik Shah, Yann LeCun, and Ramalingam Chellappa. 2022. VoLTA: Vision-language transformer with weakly-supervised local-feature alignment. ArXiv, abs/2210.04135.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.

- Deb K Roy and Alex P Pentland. 2002. [Learning words from sights and sounds: a computational model](#). *Cognitive Science*, 26(1):113–146.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. [ImageNet Large Scale Visual Recognition Challenge](#). *International Journal of Computer Vision (IJCV)*, 115(3):211–252.
- Hassan Shahmohammadi, Maria Heitmeier, Elnaz Shafaei-Bajestan, Hendrik P. A. Lensch, and Harald Baayen. 2022. [Language with vision: a study on grounded word and sentence embeddings](#).
- A. Singh, R. Hu, V. Goswami, G. Couairon, W. Galuba, M. Rohrbach, and D. Kiela. 2022. [FLAVA: A foundational language and vision alignment model](#). In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15617–15629, Los Alamitos, CA, USA. IEEE Computer Society.
- Andrew D. M. Smith and Kenny Smith. 2012. [Cross-Situational Learning](#), pages 864–866. Springer US, Boston, MA.
- Kenny Smith, Andrew D. M. Smith, and Richard A. Blythe. 2011. [Cross-Situational Learning: An Experimental Study of Word-Learning Mechanisms](#). *Cognitive Science*, 35(3):480–498.
- Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021. [WIT: Wikipedia-based image text dataset for multimodal multilingual machine learning](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 2443–2449, New York, NY, USA. Association for Computing Machinery.
- Jessica Sullivan, Michelle Mei, Andrew Perfors, Erica Wojcik, and Michael C. Frank. 2021. [SAYCam: A Large, Longitudinal Audiovisual Dataset Recorded From the Infant’s Perspective](#). *Open Mind*, 5:20–29.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: Learning cross-modality encoder representations from transformers](#). 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)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Wai Keen Vong and Brenden M. Lake. 2022. [Cross-Situational Word Learning With Multimodal Neural Networks](#). *Cognitive Science*, 46(4):e13122.
- Alex Wang and Kyunghyun Cho. 2019. [BERT has a Mouth, and It Must Speak: BERT as a Markov Random Field Language Model](#). In *Proceedings of the Workshop on Methods for Optimizing and Evaluating Neural Language Generation*, pages 30–36, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [SuperGLUE: A stickier benchmark for general-purpose language understanding systems](#). In *33rd Conference on Neural Information Processing Systems*. Journal Abbreviation: arXiv preprint arXiv:1905.00537.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Ruocheng Wang, Jiayuan Mao, Samuel Gershman, and Jiajun Wu. 2021. [Language-Mediated, Object-Centric Representation Learning](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2033–2046, Online. Association for Computational Linguistics.
- Wentao Wang, Wai Keen Vong, Najoung Kim, and Brenden M. Lake. 2023. [Finding Structure in One Child’s Linguistic Experience](#). *Cognitive Science*, 47(6):e13305.
- Alex Warstadt and Samuel R Bowman. 2022. What artificial neural networks can tell us about human language acquisition. In Shalom Lappin and Jean-Philippe Bernardy, editors, *Algebraic Structures in Natural Language*, pages 17–60. CRC Press. Publisher: CRC Press.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Gotlieb Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Adina Williams, Bhargavi Paranjabe, Tal Linzen, and Ryan Cotterell. 2023. Findings of the 2023 BabyLM Challenge: Sample-efficient pretraining on developmentally plausible corpora. In *Proceedings of the 2023 BabyLM Challenge*. Association for Computational Linguistics (ACL).
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020a. [BLiMP: The Benchmark of Linguistic Minimal Pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392. _eprint: https://doi.org/10.1162/tacl_a_00321.
- Alex Warstadt, Yian Zhang, Xiaocheng Li, Haokun Liu, and Samuel R. Bowman. 2020b. [Learning which features matter: RoBERTa acquires a preference for linguistic generalizations \(eventually\)](#). In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 217–235, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven

- Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yiqun Yao and Rada Mihalcea. 2022. [Modality-specific learning rates for effective multimodal additive late-fusion](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1824–1834, Dublin, Ireland. Association for Computational Linguistics.
- Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. 2021. [ERNIE-ViL: Knowledge Enhanced Vision-Language Representations Through Scene Graph](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(4):3208–3216.
- Tian Yun, Chen Sun, and Ellie Pavlick. 2021. [Does Vision-and-Language Pretraining Improve Lexical Grounding?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4357–4366, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yayun Zhang, Chi hsin Chen, and Chen Yu. 2019. [Chapter Two - Mechanisms of Cross-situational Learning: Behavioral and Computational Evidence](#). volume 56 of *Advances in Child Development and Behavior*, pages 37–63. JAI.
- Yian Zhang, Alex Warstadt, Xiaocheng Li, and Samuel R. Bowman. 2021. [When do you need billions of words of pretraining data?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1112–1125, Online. Association for Computational Linguistics.
- Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. 2020. [Unified Vision-Language Pre-Training for Image Captioning and VQA](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):13041–13049.

A Pretraining Validation Losses

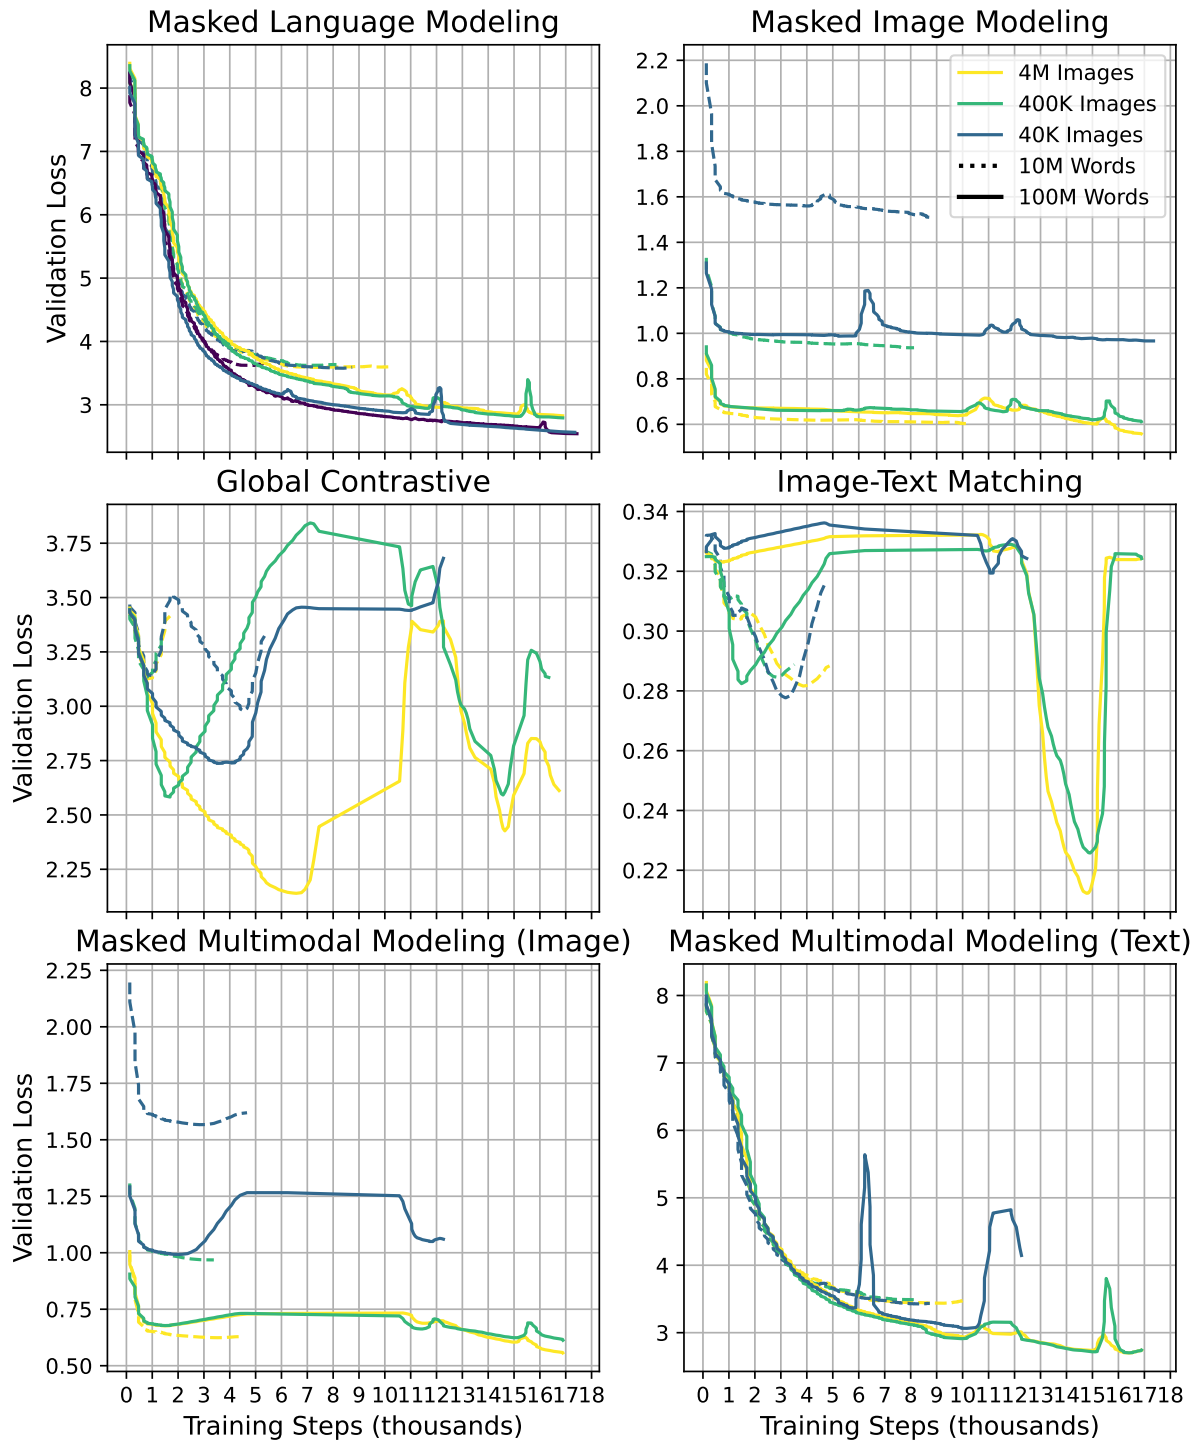


Figure 4: Validation losses for every training objective on a held-out set. While the MLM – and to a certain extent, also the MMM (Text) – losses are closely proportional to the pseudo-perplexity metric in Figure 1 (including some occasional spikes associated with checkpoint loading), the other losses are less stable. We point out some issues with the scheduler mechanism in Sections 4.1 and 4.4.

B Fine-tuning Performance

Task	Subtask	10M Words				100M Words			
		None	40K	400K	40M	None	40K	400K	40M
(Super)GLUE	BoolQ (accuracy)	64.32	65.70	65.98	<u>66.8</u>	<u>69.43</u>	67.08	65.70	66.25
	CoLA (MCC)	<u>4.21</u>	-4.16	0.00	0.00	<u>28.43</u>	21.56	20.00	20.93
	MNLI (accuracy)	<u>73.27</u>	72.46	70.07	71.98	<u>74.63</u>	71.85	74.57	74.11
	MNLI-mm (accuracy)	72.5	73.71	<u>73.98</u>	73.74	<u>78.06</u>	77.83	76.78	75.84
	MRPC (F1)	81.61	<u>83.10</u>	81.61	82.00	82.99	<u>85.51</u>	84.78	84.98
	MultiRC (accuracy)	59.58	60.35	61.45	<u>62.32</u>	<u>67.47</u>	62.76	64.62	67.25
	QNLI (accuracy)	<u>80.88</u>	79.40	80.31	79.97	<u>83.90</u>	82.76	82.55	81.98
	QQP (F1)	81.76	82.28	82.49	<u>82.57</u>	82.99	<u>84.30</u>	82.91	82.43
	RTE (accuracy)	<u>57.58</u>	53.54	53.54	55.56	53.54	57.58	<u>58.59</u>	56.57
	SST-2 (accuracy)	83.27	83.66	84.84	<u>87.01</u>	<u>91.73</u>	88.98	87.20	87.99
	WSC (accuracy)	61.45	61.45	<u>62.65</u>	61.45	<u>61.45</u>	<u>61.45</u>	60.24	<u>61.45</u>
BLiMP (Acc.)	Anaphor Agreement	76.84	<u>79.55</u>	76.99	77.04	94.89	95.35	91.10	<u>96.27</u>
	Argument Structure	<u>64.14</u>	62.38	62.55	62.45	69.58	<u>69.77</u>	68.15	68.38
	Binding	62.07	63.79	<u>64.00</u>	60.98	<u>68.91</u>	66.62	65.67	67.99
	Control/Raising	61.4	62.06	<u>64.54</u>	63.10	67.96	<u>69.64</u>	67.63	67.87
	Determiner Noun Agreement	82.91	80.42	82.11	<u>83.00</u>	90.27	<u>91.94</u>	88.45	89.49
	Ellipsis	<u>68.48</u>	65.76	61.61	66.22	<u>81.12</u>	78.12	80.60	77.37
	Filler Gap Dependencies	59.52	59.73	57.91	<u>60.85</u>	<u>67.24</u>	65.41	64.19	61.92
	Irregular Forms	<u>83.51</u>	75.98	78.73	74.45	84.17	84.12	83.66	<u>84.89</u>
	Island Effects	40.43	44.13	40.28	<u>48.02</u>	51.83	<u>53.10</u>	48.47	48.65
	NPI Licensing	52.49	54.43	<u>54.66</u>	49.53	61.24	<u>64.79</u>	60.04	54.68
	Quantifiers	58.55	57.86	<u>64.76</u>	57.73	<u>60.10</u>	58.40	57.16	58.50
Subject Verb Agreement	60.29	61.72	<u>61.81</u>	61.37	<u>75.34</u>	72.65	70.19	69.18	
MSGs (MCC)	Control Raising (control)	21.90	27.36	<u>31.42</u>	22.24	46.21	45.25	46.85	<u>55.00</u>
	Control Raising–Lexical Content	-46.54	-18.71	<u>-13.17</u>	-72.82	<u>-61.04</u>	-66.26	-63.65	-88.49
	Control Raising–Relative Position	-98.69	<u>-97.82</u>	-98.33	-100.00	-99.90	<u>-89.06</u>	-97.74	-99.34
	Lexical Content (control)	100.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00
	Main Verb (control)	86.84	<u>96.84</u>	93.24	86.92	<u>99.85</u>	99.40	96.54	98.08
	Main Verb–Lexical Content	-100.00	-100.00	-100.00	-100.00	-100.00	-100.00	-100.00	-100.00
	Main Verb–Relative Position	<u>-86.73</u>	-88.78	-91.56	-95.32	-98.41	<u>-86.56</u>	-95.64	-93.18
	Relative Position (control)	<u>96.78</u>	81.47	90.28	89.14	99.98	<u>99.47</u>	<u>100.00</u>	99.98
	Syntactic Category (control)	<u>57.04</u>	13.15	38.77	28.12	72.40	73.32	<u>77.37</u>	<u>89.29</u>
	Syntactic Category–Lexical Content	-100.00	-81.37	<u>-77.31</u>	-95.49	-81.08	<u>-74.66</u>	-79.96	-78.53
	Syntactic Category–Relative Position	-72.34	<u>-65.91</u>	-70.53	-67.61	-73.29	-68.86	-65.26	<u>-64.96</u>

Table 3: Detailed fine-tuning performance for each of the eight models. F1 denotes macro-F1, MCC denotes Matthew’s correlation coefficient, and random chance accuracy on all BLiMP tasks (i.e., the second group) is 50.