

Model Merging to Maintain Language-Only Performance in Developmentally Plausible Multimodal Models

Ece Takmaz Lisa Bylinina Jakub Dotlačil

Utrecht University

{e.k.takmaz | e.g.bylinina | j.dotlacil}@uu.nl

Abstract

State-of-the-art vision-and-language models consist of many parameters and learn from enormous datasets, surpassing the amounts of linguistic data that children are exposed to as they acquire a language. This paper presents our approach to the multimodal track of the BabyLM challenge addressing this discrepancy. We develop language-only and multimodal models in low-resource settings using developmentally plausible datasets, with our multimodal models outperforming previous BabyLM baselines. One finding in the multimodal language model literature is that these models tend to underperform in *language-only* tasks. Therefore, we focus on maintaining language-only abilities in multimodal models. To this end, we experiment with *model merging*, where we fuse the parameters of multimodal models with those of language-only models using weighted linear interpolation. Our results corroborate the findings that multimodal models underperform in language-only benchmarks that focus on grammar, and model merging with text-only models can help alleviate this problem to some extent, while maintaining multimodal performance.

1 Introduction

Current state-of-the-art multimodal language models (MLMs) are composed of many layers containing billions of parameters and they require huge amounts of data to learn how to handle and bridge visual and textual modalities. On the other hand, children acquire language with the help of much smaller sets of linguistic input. The BabyLM challenge (Warstadt et al., 2023) focuses on this discrepancy and encourages the implementation and training of sample-efficient, developmentally plausible models in resource-limited contexts. Although utilizing small datasets and models could prove challenging to outperform current MLMs, such setups could allow for cognitive plausibility, also making the development and use of such models more




Figure 1: Weighted merging of language-only and multimodal models in the form of linear interpolation.

accessible and efficient.

Despite the general good performance of MLMs in multimodal tasks, previous work shows that these models tend to underperform in *language-only* tasks (Zhuang et al., 2024). Recent work in the multimodal BabyLM challenge also points to the same issue (Amariuca and Warstadt, 2023; Klerings et al., 2024). Therefore, our aim in this paper is to first test our own models on multimodal and text-only benchmarks, and second, if we observe the same issue, to try to mitigate it.

We develop language-only and multimodal models, the latter of which outperforms previous BabyLM baselines.¹ However, our results indeed confirm that our developmentally plausible MLMs lack in text-only benchmarks. Hence, we explore a model augmentation technique to potentially overcome this shortcoming: **model merging**. Model merging has been utilized to prevent catastrophic forgetting and combine the capabilities of multiple models trained on different tasks, datasets, or modalities (Yang et al., 2024; Dash et al., 2025).

In our approach, during inference time, we fuse the parameters of models trained on text-only and multimodal data in a straightforward, training-free way (see Figure 1). Our results indicate that such an augmentation yields a single model maintaining accuracy and robustness in both text-only and multimodal benchmarks in the earlier and later stages of training the multimodal model.

¹Code and models available at https://github.com/ecekt/babylm_multimodal_model_merging

2 Background

We first go into detail about the multimodal approaches to the BabyLM challenge in Section 2.1, and then, provide a summary of the recent work on model merging in Section 2.2.

2.1 Developmentally Plausible Multimodal Language Models

The BabyLM initiative encourages the development of models that can be small-scale, trained on smaller sets of data, using various techniques such as model compression, learning from interaction, knowledge distillation. Our focus is on the **multimodal track**. Multimodal models have been explored in the 2nd BabyLM challenge (Choshen et al., 2024) and re-introduced in the 3rd BabyLM challenge (Charpentier et al., 2025) as the submitted models did not outperform the baselines released by the BabyLM organizers (Hu et al., 2024). The baselines were GIT (Wang et al., 2022) and Flamingo (Alayrac et al., 2022) models trained on the BabyLM’s multimodal corpus to ground language to vision.

These models encode image inputs and generate text using text decoders conditioned on visual tokens. The methodologies from the past submissions include curriculum learning where the captions were ordered based on the number of concepts they included (Saha et al., 2024), where this helped on developmentally plausible benchmarks. Pre-training on text also seems to be beneficial; Saha et al. (2024) investigates first training on text and then captions along with curriculum learning using image-caption pairs. However, in general, it appears to be difficult to observe a strong pattern across model types, datasets and tasks.

Another work reports a related result where learning in phases appears to benefit multimodal BabyLM models (AlKhamissi et al., 2024). The model first learns the language-only tasks, then grounding, followed by self-synthesized data and more advanced reasoning tasks.

There are also contributions to the language-only track where the models were influenced or informed by multimodal input (Fields et al., 2023; Amariuca and Warstadt, 2023).

Klerings et al. (2024) explore a weighted loss function for text-only and multimodal data during training. However, they show that vision does not significantly benefit the performance in language-only benchmarks. This is in line with prior findings

showing limited or no improvements when incorporating visual data (Amariuca and Warstadt, 2023; Zhuang et al., 2024), with the exception of low-data regimes (Zhuang et al., 2024), which inspired our work.

2.2 Model Merging

Model merging has been utilized as a technique for adaptively extending the capabilities of models or balancing performance during inference time in the tasks multiple models were trained on. See (Yang et al., 2024) and (Goddard et al., 2024) for surveys of various merging techniques.

A straightforward averaging technique called ‘model soups’ has been found beneficial in improving accuracy and robustness. The techniques involve combining the parameters of multiple models trained on different hyperparameters, in addition to more sophisticated weighted merging methods (Wortsman et al., 2022; Matena and Raffel, 2022). Similarly, Aakanksha et al. (2024) find that merging models is better than mixing training data for facilitating safety and multilingual generalizability.

Regarding vision-and-language models, Zhu et al. (2025) learn modules for various multimodal tasks that are later merged; whereas (Li et al., 2025) exploit text-only reward models to transfer to vision-and-language reward models in a cross-modal model merging scheme.

Closer to our approach, AyaVision is an example of cross-modal merging to maintain text-only capabilities within multimodal models to prevent catastrophic forgetting (Dash et al., 2025). The authors built their multimodal model on their best-performing text-only checkpoint, which makes the setup more suitable for merging. Similarly, Sung et al. (2023) conduct detailed experiments on multimodal model merging, finding that simple linear interpolation is a competitive and efficient method, which we also opt for in this work to test its effectiveness in low-data and low-compute settings.

3 Data

To train our models, we use the data from the multimodal BabyLM challenge, which consists of 2 parts: text-only and multimodal.

Text-only. We use the 50M-word text data provided by the BabyLM challenge. This data consists of text stemming from 6 sources as explained by Choshen et al. (2024), and corresponds to the

Data	Train	Val
Localized Narratives	729349	38387
CC3M	2061837	108518
BabyLM - text-only	5492930	289102
Total	8284116	436007

Table 1: Number of samples per dataset in the splits we created (after filtering, validating and deduplicating CC3M images and captions, and trimming it to fit 100M words in total).

first halves of the text-only subsets released for the language-only challenge of the BabyLM task.

Multimodal. This part includes image-caption pairs from Localized Narratives (Pont-Tuset et al., 2020) and Conceptual Captions 3M (Sharma et al., 2018). We download the Localized Narratives (LN) images and captions from the dataset’s website.² We use the COCO (train) (Lin et al., 2014) and Open Images (train and test) (Kuznetsova et al., 2020) subsets of LN.³

Additionally, we download the captions for the existing images from the Conceptual Captions (CC3M) dataset.⁴ Filtering out the images that do not exist anymore as well as the corrupt and duplicate image files, we end up with fewer than 3M images, which is lower than the provided captions for the previous multimodal BabyLM challenge.

Final dataset. The statistics of our final dataset are provided in Table 1, with our random train and validation splits where 95% of each subset contributes to the training set and the rest goes into the validation set.

4 Methodology

Models. We modify the implementation of LLaVa (Liu et al., 2023, 2024) from HuggingFace⁵, inheriting the LlavaForConditionalGeneration model, and replace the visual encoder with

²<https://google.github.io/localized-narratives/>

³Although the BabyLM OSF repository at <https://osf.io/ad7qg/files> provides captions and extracted image representations for this subset, we noticed a discrepancy in the number of samples compared to the original LN. It seems that the test set of the Open Images subset is also counted in the BabyLM corpus to end up with the 50M word count. Therefore, we also included that part. Localized Narrative subset IDs of Open Image downloaded from https://storage.googleapis.com/openimages/web/download_v7.html. COCO images from <https://cocodataset.org/>

⁴We download the captions from <https://ai.google.com/research/ConceptualCaptions/download> and the images using the script provided at <https://github.com/igorbrigadir/DownloadConceptualCaptions>.

⁵llava-hf/llava-1.5-7b-hf’

DINOv2-large (Oquab et al., 2024)⁶. We also make necessary changes to the image processing code and modeling code in relation to the dimensions of the image features. We randomly initialize a 6-layer version of this model as the language model, together with a mapping layer that projects the image representations to the language model’s space. **Image representations.** Unlike the BabyLM benchmark’s image representations from last year that are 768-dimensional vectors from DINOv2 ViT-Base, we use the large version of DINOv2, which processes images into 256 image tokens of 1024 dimensions. While we originally intended to feed all 256 image tokens extracted from the vision encoder, due to time and compute constraints, we modified the model to feed a single pooled image token directly. We implemented a version of the model where we pre-extract all image token representations and mean-pool them. This single summary token (1024 dim) is fed to the LM directly (bypassing the vision tower). The summary image representation goes through a multimodal projector composed of a linear layer projecting from 1024 to 768 dimensions, GeLU activation and another linear layer projecting from 768 to 768. This multimodal projector is trained along with the language model, while keeping the image representations frozen.

For text-only data, we create a black image (640 x 420) and always input the features of this placeholder image both in the text-only model and the multimodal model.

Training the tokenizer. Using all the text in the final dataset (token count = 100M), we train a tokenizer from scratch employing the configurations of the LLaVA tokenizer (LlamaTokenizerFast, a byte-pair encoding model based on SentencePiece), with a vocabulary size of 30000 including a special token for image representations. Using the BERT pre-tokenizer, we apply splitting on whitespace and punctuation. This preprocessing yields a 1.36 word-to-subword ratio.⁷

Intermediate checkpoints. To investigate the learning speed and model behavior dynamics, we save checkpoints (every 1M words up to 10M, every 10M up to 100M, every 100M up to 1B). We estimate the words-seen using the ratio of word-pieces to actual words in our dataset (1.36). We use this ratio to roughly determine how many ‘words’

⁶facebook/dinov2-large’

⁷Words: 99,999,990. Subwords (as tokenized by our tokenizer, skipping special tokens): 136,034,832.

the models were exposed to in the training batches (excluding special tokens).

Hyperparameters and setup. We follow the restrictions of the BabyLM challenge, 100M words, 10 epochs, resulting in 1B tokens seen in total.

We set the maximum length to 150 tokens. We truncate longer samples if they do not fit this constraint; if they are shorter, we pad them.

We train the models on 2 A10 GPUs with 24GB memory on CrossEntropy Loss using the AdamW optimizer with a learning rate of $1e-4$. For the multimodal model, batches can contain text-only or multimodal data, and the loss is calculated in the same way for both modalities. We use fp16 half precision and make use of the Accelerate library for data parallelism to speed up training. We accumulate gradients for 8 steps and then apply gradient updates to optimize the model, effectively increasing our batch sizes from 64 to 512. The numbers of words seen are gathered from each GPU and only logged and checked in the main process. We opted for a smaller layer number (6) to allow for a speed-up in the training by exposing the model to larger and more batches in a shorter amount of time. It takes 6.7 days for the multimodal model (with text-only and multimodal data) to be trained, and the text-only model 4.3 days.

5 Model Merging at Inference Time

Since our language-only and multimodal models share the same architecture, random initialization and the text-only data, they can be combined in a straightforward way. We apply a simple weighted sum of the multimodal model’s parameters and the text-only model’s parameters. We experiment with merging weights α of 0.3, 0.5, and 0.8,⁸ where θ indicates all the trainable parameters of a model:

$$\theta_{merged} = \alpha\theta_{LLM} + (1 - \alpha)\theta_{VLM} \quad (1)$$

In this way, the merged model is a linear interpolation of the multimodal and text-only models.

6 Benchmarks

We modify the evaluation pipeline provided by the BabyLM challenge⁹ to run zero-shot evaluations

⁸The weights were chosen to reflect equal contribution from both models (0.5) and a skewed contribution from one model (0.3—more VLM and 0.8—more LLM).

⁹<https://github.com/babylm/evaluation-pipeline-2025>

across our checkpoints using the benchmarks provided. For Winoground, we write our own evaluation code.

Language-only benchmarks. We run the evaluation pipeline for all the tasks in BLiMP (Warstadt et al., 2020), EWoK, entity tracking (assign the highest probability to the correct continuation) (Kim and Schuster, 2023), Wug past tense (Weissweiler et al., 2023), wug adjective nominalization (Hofmann et al., 2025) testing morphological capabilities (correlating model probabilities to human judgments). BLiMP and BLiMP-supplement (more challenging samples) evaluate whether the models capture grammatical phenomena, where one grammatical and one ungrammatical sentence are pitted against each other, testing models’ capabilities related to syntax, morphology and semantics. EWoK (Ivanova et al., 2025) focuses on world knowledge and reasoning about e.g., social, physical, spatial relations.

Multimodal benchmark. We experiment on Winoground (Thrush et al., 2022), where pairs of images with very similar captions are provided. Winoground consists of 400 samples, where each sample has 2 images and 2 captions. These 2 captions have the same words, but in different orders to match the image contents (see an example in Figure 2).



Figure 2: An example from Winoground. *Left*: ‘painting the white wall red’. *Right*: ‘painting the red wall white’.

We input the image and 2 captions separately to the model to obtain predictions for Winoground. If the likelihood of the correct caption is higher, we increase the accuracy. We use the unpaired text-score as used in previous BabyLM work, where we consider each image-caption pair separately. We use the full Winoground dataset available on HuggingFace, unlike the filtered version in the BabyLM evaluation suite.

Winoground tests abilities requiring compositionality, sensitivity to word order, common-sense reasoning, pragmatics and overall more fine-grained visual and linguistic analyses involving unusual images and texts (Diwan et al., 2022).




Figure 3: Average accuracies for the text-only model and the multimodal model over the training checkpoints, for the BLiMP full benchmark.




Figure 4: Average accuracies for the text-only model and the multimodal model over the training checkpoints, for the BLiMP fast benchmark.

Winoground is a difficult dataset, with previous BabyLM work yielding accuracies as follows: 2024 baselines Flamingo: 51.6 GIT: 55.5; 2025 baselines Flamingo: 54.8, GIT, 56.2.¹⁰

7 Results

Results on benchmarks. We first obtain the results on the full BLiMP evaluation, which is reported in Figure 3. Our best language-only model reaches 74.82 accuracy at the 100M checkpoint. **Our best multimodal model yields 73.84 accuracy, surpassing the multimodal BabyLM baselines as well as the current 2 submissions on the multimodal BabyLM leaderboard** (2024 Flamingo: 70.9, GIT: 65.2, 2025 Flamingo: 70.9, GIT: 72.2). We see that, generally, the text-only model outperforms or is on par with the multimodal model, except some later checkpoints.

We use the ‘fast’ versions of the benchmarks that contain a smaller set of samples to obtain the following results due to time and compute constraints.¹¹ In Figure 4, we depict the performance

¹⁰2024 baselines from: <https://github.com/babylm/evaluation-pipeline-2024/>, 2025 baselines from: <https://huggingface.co/spaces/BabyLM-community/babylm-leaderboard-2025-all-tasks>

¹¹We noticed that the Wug fast and full benchmarks are in

of the text-only and multimodal model checkpoints on the BLiMP fast benchmark, which yields outcomes closely aligned with those obtained from the full benchmark.

We see similar trends for BLiMP supplement (Figure 5), Wug past tense (Figure 6) and adjective nominalization (Figure 7) benchmarks. This is in line with previous work indicating that multimodal data tend not to benefit performance on language-only benchmarks.

When we look at the results on the Entity Tracking (Figure 8) and EWoK benchmarks (Figure 9), however, we see trends where multimodal checkpoints clearly outperform the text-only checkpoints. This could be due to the focus of these datasets, which is more knowledge- and semantics-oriented rather than grammatical, therefore, the multimodal data such as the image descriptions in narrative form from the Localized Narratives dataset could have helped.

Although our models do not perform well in the BLiMP supplement and Wug past tense benchmarks, they show competitive performance in the remaining tasks.

Results on model merging. We use BLiMP as a

fact identical.




Figure 5: Average accuracies for the text-only model and the multimodal model over the training checkpoints, for the BLiMP supplement fast benchmark.




Figure 6: Correlation between model predictions and human responses from the Wug past tense benchmark, for the text-only model and the multimodal model over the training checkpoints.





Figure 7: Correlation between model predictions and human responses from the Wug adjective nominalization benchmark, for the text-only model and the multimodal model over the training checkpoints.






Figure 9: Average accuracies for the text-only model and the multimodal model over the training checkpoints, for the EWOK fast benchmark.




Figure 10: Average accuracies for the merged models with different weights (higher α indicates more contributions from the language-only model), along with the training dynamics of the multimodal model for the BLiMP fast benchmark.




Figure 11: Average accuracies for Winoground. MM represents the multimodal model checkpoints and MM Merged indicates a merged model with $\alpha = 0.3$, using the language-only checkpoint with the highest BLiMP score. The red dotted line indicates random chance accuracy. The top score in the leaderboard on the filtered version is 56.2.

use case for our model merging experiments. We merge each multimodal checkpoint model with the language-only model that performs best in BLiMP (100M checkpoint) with varying weights. Figure 10 illustrates the model-merging results when combining the language-only and multimodal models using $\alpha = 0.3, 0.5$ and 0.8 as the weights of the language-only model and $1 - \alpha$ for the multimodal model. Merging the trained language-only model in the early training stages of the multimodal model meaningfully helps in getting better results in BLiMP. Additionally, in the later checkpoints when the multimodal model’s language-only capabilities begin to drop, 0.3/0.7 merging scheme helps the model maintain language-only capabilities.

To check whether merging with the language-only model affects multimodal performance, we also look at the accuracy on the Winoground benchmark after merging models. The results for Winoground are provided in Figure 11, showing that in some checkpoints, merging can actually be beneficial without significantly decreasing multimodal scores.

8 Conclusion

We have investigated whether model modification in the form of model merging at inference time would benefit multimodal BabyLM models in language-only and multimodal tasks. Our results showed that, indeed, multimodal models tend to underperform in text-only benchmarks that focus on grammar (although surpassing previous baselines) and model merging with text-only models can help alleviate this issue to some extent. Future work can explore other model merging techniques and the effects of model merging in other benchmarks.

Limitations

Due to time and compute constraints, we altered our intended initial setup where the model is fed 256 image patches into one where a single, pooled image representation is relayed to the model. This might cause information loss and performance drop, and ideally, we would like to provide the whole set of image patches. We tested models with 6 transformer layers, which is quite few compared to state-of-the-art models. Therefore, this might have resulted in lesser performance. However, we believe that our results shed light on what to expect in compute and data-efficient/scarcely setups, which should be investigated further using more

seeds and different training orders for robustness and generalizability of the conclusions. Additionally, although the set of benchmarks we tested on does not cover the full spectrum of language-only and multimodal tasks in the BabyLM challenge, we think that they span a reasonable range of them, providing insights into the dynamics visuo-linguistic processes as training progresses.

Acknowledgments

The research reported in this paper was supported by the European Research Council (ERC), grant 101088098 - MEMLANG. The first author utilized compute resources granted by the Dutch Research Council (NWO). The second author acknowledges support from NWO, Open Competition XS project number 406.XS.25.01.104. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Aakanksha, Arash Ahmadian, Seraphina Goldfarb-Tarrant, Beyza Ermis, Marzieh Fadaee, and Sara Hooker. 2024. [Mix data or merge models? optimizing for performance and safety in multilingual contexts](#). In *Neurips Safe Generative AI Workshop 2024*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, and 8 others. 2022. [Flamingo: a visual language model for few-shot learning](#). In *Advances in Neural Information Processing Systems*.
- Badr AlKhamissi, Yingtian Tang, Abdülkadir Gökce, Johannes Mehrer, and Martin Schrimpf. 2024. [Dreaming out loud: A self-synthesis approach for training vision-language models with developmentally plausible data](#). In *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pages 244–251, Miami, FL, USA. Association for Computational Linguistics.
- Theodor Amariuca and Alexander Scott Warstadt. 2023. [Acquiring linguistic knowledge from multimodal input](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 128–141, Singapore. Association for Computational Linguistics.

- Lucas Charpentier, Leshem Choshen, Ryan Cotterell, Mustafa Omer Gul, Michael Hu, Jaap Jumelet, Tal Linzen, Jing Liu, Aaron Mueller, Candace Ross, Raj Sanjay Shah, Alex Warstadt, Ethan Wilcox, and Adina Williams. 2025. [BabyLM turns 3: Call for papers for the 2025 babyLM workshop](#). *Preprint*, arXiv:2502.10645.
- Leshem Choshen, Ryan Cotterell, Michael Y. Hu, Tal Linzen, Aaron Mueller, Candace Ross, Alex Warstadt, Ethan Wilcox, Adina Williams, and Chengxu Zhuang. 2024. [\[call for papers\] the 2nd babyLM challenge: Sample-efficient pretraining on a developmentally plausible corpus](#). *Preprint*, arXiv:2404.06214.
- Saurabh Dash, Yiyang Nan, John Dang, Arash Ahmadian, Shivalika Singh, Madeline Smith, Bharat Venkitesh, Vlad Shmyhlo, Viraat Aryabumi, Walter Beller-Morales, Jeremy Pekmez, Jason Ozuzu, Pierre Richemond, Acyr Locatelli, Nick Frosst, Phil Blunsom, Aidan Gomez, Ivan Zhang, Marzieh Fadaee, and 6 others. 2025. [Aya vision: Advancing the frontier of multilingual multimodality](#). *Preprint*, arXiv:2505.08751.
- Anuj Diwan, Layne Berry, Eunsol Choi, David Harwath, and Kyle Mahowald. 2022. [Why is winoground hard? investigating failures in visuo-linguistic compositionality](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2236–2250, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Clayton Fields, Osama Natouf, Andrew McMains, Catherine Henry, and Casey Kennington. 2023. [Tiny language models enriched with multimodal knowledge from multiplex networks](#). In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 47–57, Singapore. Association for Computational Linguistics.
- Charles Goddard, Shamane Siriwardhana, Malikeh Ehghaghi, Luke Meyers, Vladimir Karpukhin, Brian Benedict, Mark McQuade, and Jacob Solawetz. 2024. [Arcee’s MergeKit: A toolkit for merging large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 477–485, Miami, Florida, US. Association for Computational Linguistics.
- Valentin Hofmann, Leonie Weissweiler, David R. Mortensen, Hinrich Schütze, and Janet B. Pierrehumbert. 2025. [Derivational morphology reveals analogical generalization in large language models](#). *Proceedings of the National Academy of Sciences*, 122(19):e2423232122.
- Michael Y. Hu, Aaron Mueller, Candace Ross, Adina Williams, Tal Linzen, Chengxu Zhuang, Ryan Cotterell, Leshem Choshen, Alex Warstadt, and Ethan Gottlieb Wilcox. 2024. [Findings of the second BabyLM challenge: Sample-efficient pretraining on developmentally plausible corpora](#). In *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pages 1–21, Miami, FL, USA. Association for Computational Linguistics.
- Anna A. Ivanova, Aalok Sathe, Benjamin Lipkin, Unnathi Kumar, Setayesh Radkani, Thomas H. Clark, Carina Kauf, Jennifer Hu, R. T. Pramod, Gabriel Grand, Vivian Paulun, Maria Ryskina, Ekin Akyürek, Ethan Wilcox, Nafisa Rashid, Leshem Choshen, Roger Levy, Evelina Fedorenko, Joshua Tenenbaum, and Jacob Andreas. 2025. [Elements of world knowledge \(ewok\): A cognition-inspired framework for evaluating basic world knowledge in language models](#). *Preprint*, arXiv:2405.09605.
- Najoung Kim and Sebastian Schuster. 2023. [Entity tracking in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3835–3855, Toronto, Canada. Association for Computational Linguistics.
- Alina Klerings, Christian Bartelt, and Aaron Mueller. 2024. [Developmentally plausible multimodal language models are highly modular](#). In *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pages 118–139, Miami, FL, USA. Association for Computational Linguistics.
- Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper R. R. Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. 2020. [The open images dataset V4](#). *International journal of computer vision*, 128(7):1956–1981.
- Chen-An Li, Tzu-Han Lin, Yun-Nung Chen, and Hung-yi Lee. 2025. [Transferring textual preferences to vision-language understanding through model merging](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 923–943, Vienna, Austria. Association for Computational Linguistics.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. [Microsoft COCO: Common objects in context](#). In *Computer Vision – ECCV 2014*, pages 740–755, Cham. Springer International Publishing.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. [Improved baselines with visual instruction tuning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.

- Michael Matena and Colin Raffel. 2022. Merging models with fisher-weighted averaging. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, and 7 others. 2024. **DINOv2: Learning robust visual features without supervision**. *Transactions on Machine Learning Research*. Featured Certification.
- Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. 2020. Connecting vision and language with localized narratives. In *ECCV*.
- Rohan Saha, Abrar Fahim, Alona Fyshe, and Alex Murphy. 2024. **Exploring curriculum learning for vision-language tasks: A study on small-scale multimodal training**. In *The 2nd BabyLM Challenge at the 28th Conference on Computational Natural Language Learning*, pages 65–81, Miami, FL, USA. Association for Computational Linguistics.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. **Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, Melbourne, Australia. Association for Computational Linguistics.
- Yi-Lin Sung, Linjie Li, Kevin Lin, Zhe Gan, Mohit Bansal, and Lijuan Wang. 2023. **An empirical study of multimodal model merging**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1563–1575, Singapore. Association for Computational Linguistics.
- Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. 2022. Winoground: Probing vision and language models for visio-linguistic compositionality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5238–5248.
- Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. 2022. **GIT: A generative image-to-text transformer for vision and language**. *Transactions on Machine Learning Research*.
- Alex Warstadt, Aaron Mueller, Leshem Choshen, Ethan Wilcox, Chengxu Zhuang, Juan Ciro, Rafael Mosquera, Bhargavi Paranjabe, Adina Williams, Tal Linzen, and Ryan Cotterell. 2023. **Findings of the BabyLM challenge: Sample-efficient pretraining on developmentally plausible corpora**. In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 1–34, Singapore. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. **Blimp: The benchmark of linguistic minimal pairs for english**. *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Leonie Weissweiler, Valentin Hofmann, Anjali Kantharuban, Anna Cai, Ritam Dutt, Amey Hengle, Anubha Kabra, Atharva Kulkarni, Abhishek Vijayakumar, Haofei Yu, Hinrich Schuetze, Kemal Oflazer, and David Mortensen. 2023. **Counting the bugs in ChatGPT’s wugs: A multilingual investigation into the morphological capabilities of a large language model**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6508–6524, Singapore. Association for Computational Linguistics.
- Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and 1 others. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *International conference on machine learning*, pages 23965–23998. PMLR.
- Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang, Xiaochun Cao, Jie Zhang, and Dacheng Tao. 2024. **Model merging in llms, mllms, and beyond: Methods, theories, applications and opportunities**. *Preprint*, arXiv:2408.07666.
- Didi Zhu, Yibing Song, Tao Shen, Ziyu Zhao, Jinluan Yang, Min Zhang, and Chao Wu. 2025. **REMEDY: Recipe merging dynamics in large vision-language models**. In *The Thirteenth International Conference on Learning Representations*.
- Chengxu Zhuang, Evelina Fedorenko, and Jacob Andreas. 2024. **Visual grounding helps learn word meanings in low-data regimes**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1311–1329, Mexico City, Mexico. Association for Computational Linguistics.