

Analyzing the Roles of Language and Vision in Learning from Limited Data

Allison Chen¹, Ilia Sucholutsky¹, Olga Russakovsky¹, Thomas L. Griffiths^{1,2}

¹ Department of Computer Science, Princeton University

² Department of Psychology, Princeton University

{allisonchen, is2961, olgarus, tomg}@princeton.edu

Abstract

Does language help make sense of the visual world? How important is it to actually see the world rather than having it described with words? These basic questions about the nature of intelligence have been difficult to answer because we only had one example of an intelligent system – humans – and limited access to cases that isolated language or vision. However, the development of sophisticated Vision-Language Models (VLMs) by artificial intelligence researchers offers us new opportunities to explore the contributions that language and vision make to learning about the world. We ablate components from the cognitive architecture of these models to identify their contributions to learning new tasks from limited data. We find that a language model leveraging all components recovers a majority of a VLM’s performance, despite its lack of visual input, and that language seems to allow this by providing access to prior knowledge and reasoning.

Keywords: language; visual understanding; vision-language; cognitive architectures

Introduction

Language is effective in communicating visual ideas between two individuals, but what is its role in *enabling* visual understanding and intelligence? Perhaps we might attribute the effectiveness of language in conveying visual ideas to its combination with visual input through reasoning. But is it possible that language itself can facilitate visual understanding? If so, to what extent? Intuitively, we might think that seeing the world is necessary for understanding its visual structure, but studies show that congenitally blind (i.e., blind since birth) individuals have mental representations of color similar to those of sighted individuals (Kim et al., 2021). Similar results have been found for other visual tasks involving spatial memorization, shape understanding, and visual concept association (Kerr, 1983; Aleman et al., 2001; Zimler & Keenan, 1983; Connolly et al., 2007).

Until now, studies like these that investigate where nature has isolated language or vision have offered our best chance at identifying the roles that these different capacities have in our understanding of the world around us. However, recent developments in artificial intelligence offer another way to tease apart the contributions of language and vision. They provide tools to create new kinds of intelligent systems, enabling us to explore the implications of using cognitive architectures that differ in inputs and components.

Multi-modal Vision-Language Models (VLMs) (Radford et al., 2021; Liu et al., 2023; Singh et al., 2022) are ideal ex-

amples of systems that can be used to engage with these classic questions. VLMs are built by combining vision models with Large Language Models (LLMs) and have demonstrated impressive visual recognition and understanding (X. Sun et al., 2022; Guo et al., 2023; Rao et al., 2022). By studying a VLM and its corresponding LLM backbone, we can isolate vision and language capabilities to analyze their respective contribution to visual understanding. By simulating systems with different components, we can explore how the cognitive architectures of these systems affect their performance.

The key contribution of our work is understanding how each component in a VLM, especially the language-related ones, contributes to its performance. We focus on three internal components: visual processing, prior knowledge, and reasoning. We start with a VLM, break down its cognitive architecture using related models, and analyze each component’s contribution to performance on learning new tasks from limited data. Our results show that removing visual processing but retaining full language capabilities maintains just over 75% of the VLM performance. Furthermore, when any one of the language components is missing, performance drops significantly, indicating each one is necessary. Conversely, reducing a VLM to a vision-only model cuts performance in half and is comparable to an incomplete language model. These results suggest that language may indeed play a key role in understanding the visual world, allowing us to use prior knowledge and reasoning to make sense of our sensory information.

Background

In this section, we will first review LLMs and VLMs then summarize results in cognitive science that support the idea that language may enable visual understanding.

Large Language Models & Vision-Language Models

LLMs are large computational models, typically possessing billions of parameters, that use massive amounts of text data to optimize the next word prediction in a given sequence. Common models include the GPT family (Radford et al., 2018, 2019; Brown et al., 2020; OpenAI, 2023), LLaMA (Touvron et al., 2023), and Vicuna (Zheng et al., 2023). These models use the Transformer architecture (Vaswani et al., 2017) to learn which previous words to attend to when predicting the next word. Much of an LLM’s cognitive abil-

ities can be attributed to one of two components: reasoning and prior knowledge. Reasoning provides the ability to adapt to new tasks via in-context learning, i.e., providing examples in the prompt (Brown et al., 2020; Gao et al., 2021; Bansal et al., 2020). Prior knowledge enables LLMs to answer factual questions (Petroni et al., 2019; Heinzerling & Inui, 2021). In our work, we aim to isolate these two components by varying the task at hand to understand how they each contribute to performing visual recognition.

VLMs connect LLMs to vision models. Early VLMs such as CLIP (Radford et al., 2021), ALIGN (Jia et al., 2021), and DeCLIP (Li et al., 2022) aligned the representation spaces between text and image encoders using methods such as contrastive learning (Chopra et al., 2005; Hadsell et al., 2006). However, as models grew in size, training them became computationally prohibitive. As a result, newer VLMs such as LLaVA (Liu et al., 2023) adapt pre-trained image encoders (Dosovitskiy et al., 2021; Radford et al., 2021) to fit the representation space of a pre-trained LLM. VLMs have been applied to video understanding (Wu et al., 2023; Zhao et al., 2023), robotic planning (Ahn et al., 2022; Driess et al., 2023), and even geometry-based tasks such as pose estimation (Feng et al., 2023). This widespread adoption suggests that language enhances a vision model’s capabilities. However, by entangling the two, it is difficult to pinpoint what language specifically contributes to visual understanding.

Language and Visual Understanding

It may seem counterintuitive that without visual inputs, language alone could come close to understanding images. However, as mentioned earlier, congenitally blind individuals demonstrate a similar understanding of many visual concepts to sighted individuals (Kim et al., 2021; Kerr, 1983; Aleman et al., 2001; Zimler & Keenan, 1983; Connolly et al., 2007). This suggests that there is potential to use linguistic descriptions to build accurate representations of visual input.

Additionally, recent work suggests that purely text-based LLMs create representations that correspond well with those of humans in settings that range from sensory domains such as taste, color, and timbre (Marjeh et al., 2023), image/video/audio representations (Marjeh, Sucholutsky, et al., 2022; Marjeh, Van Rijn, et al., 2022), and even abstract concepts (Xu et al., 2023). These findings suggest that text is sufficient to give LLMs an understanding of the basic sensory inputs, visual similarity, and abstract concept associations that underlie our representation of the world. Based on these results, we hypothesize that language may play a substantial role in visual understanding; we extend these works by moving beyond representational similarity and into visual understanding to explore this hypothesis.

Exploring the Space of Cognitive Architectures

The question of what roles language and vision play in learning from limited data about the visual world around us can be framed as a question about *cognitive architectures*. Cognitive

architectures provide both a theoretical structure for understanding the components of complex cognitive systems and a practical framework for implementing computational models. First introduced in Newell et al. (1972), they have since been used as the basis for models of human cognition, such as in Anderson (2013). More recently, researchers have looked to apply cognitive architectures to systems based upon LLMs (Sumers et al., 2023; R. Sun, 2024) as a way to better understand, design, and compare the assumptions in these systems.

We can use the cognitive architecture approach to hypothesize the *abstract* components that contribute to learning to solve a new visual recognition problem. Fig. 1 shows a simple architecture for accomplishing this. It identifies four important components of learning visual understanding from limited data, three of which are internal to the model: **visual processing**, prior **knowledge**, and **reasoning**. The fourth is relevant training **examples**. Both humans and VLMs can be analyzed in terms of this architecture. In VLMs, the image model provides visual processing, knowledge is the stored internal representations of both the vision and language models, and reasoning is the non-trivial method of piecing information together to form a final prediction.




Figure 1: Hypothesized cognitive architecture for intelligent visual recognition systems. A full architecture consists of visual processing, prior knowledge, and reasoning components and is given relevant training examples prior to testing.

Having described VLMs and LLMs in these terms, we can then effectively ablate specific components of these models to determine their significance. Concretely, we differentiate between LLMs and VLMs depending on the input type: text only for LLMs and text and images for VLMs.

We use the task of image recognition to measure visual understanding from limited data. We begin with a full VLM containing all three components of the architecture (i.e., “Vision + Knowledge + Reasoning”). To study language’s abilities to enable visual understanding, we isolate the LLM (abstaining from visual processing), use semantically meaningful text representations of each image (rather than the pixel-based image), and ablate components of the model. We also remove language from the VLM by utilizing a vision-only model with images and perform similar ablations. Removing each component and re-testing the performance of the system illustrates the contribution and necessity of each component.

Our simulations address three specific research questions.





	Image	Tags	Caption	Label
Clean		crab, aquarium, saltwater, stone crab, bad crab	we're not exactly sure what kind of crab it is. it showed up one day in sue's salt-water tank. probably a hitchhiker on a rock with corals. it was rather shy[...]	crab
Typos		tissage, jewelry, handmade, necklace, aqua, faceted, teal, flower, light blue, formal, speciall occasion, chunky, romantic, silver	the focal piece is a rhinestone and enamel beauty with a lucite cabochon. for the rest of the necklace i used gorgeous faceted aqua blue glass beads[...]	necklace
Non-descriptive caption		cavy exhibition, cavy show, guinea pig show, guinea pig, guinea, pig, cavy, cavies, pet, cute, adorable, rodent, pocket pet, christmas show, xmas show	junior exhibitor class.	guinea pig
Non-descriptive tags and captions		25+faves	larger?	pepper

Figure 2: Examples of images, corresponding tags, captions, and the class labels. [...] indicates the text was truncated due to length. First row represents a clean example and the bottom three demonstrate variance in the data.

First, how close can an LLM (lacking visual processing) get to the performance of a VLM (with visual processing)? Second, how do the LLM’s various components contribute to its performance? And third, how does a vision model lacking language compare to the VLM?

Methods

We focus on the task of image recognition to measure visual understanding. In this section, we explain our dataset curation process, describe our instantiations of each cognitive architecture, and outline the three simulations we conducted.

Data Curation

We sought a large dataset with semantically meaningful text-image pairs with associated class labels from a finite set of classes. We chose ImageNet Captions (Fang et al., 2022), which is based on the ImageNet Large Scale Visual Recognition Challenge 2012 (Russakovsky et al., 2015) training set. Each image is associated with an ImageNet-1K class label as well as a title, caption, and/or tags from Flickr, where the images were originally uploaded. As visualized in Fig. 2, some of the tags and captions are non-informative. However, this variability enables us to probe the limits of accurate classification while facing limits on descriptiveness due to realistic noise. Further, we used baselines and compared results across consistent test sets such that the effect of uninformative text is equivalent across conditions. For our simulations, we only used the tags as text representations due to their high information density, and we generally found that they are sufficiently descriptive of the images to make classification possible.

Preprocessing included cleaning and filtering the data and selecting a subset to use for classification. The ImageNet

Captions dataset originally contained 463.6K images spanning 999 of the ImageNet-1K classes with between 7 to 961 samples per class. Text preprocessing entailed replacing all continuous white space with a single space, removing HTML tags, and converting all text to lowercase. Next, we only kept samples with all three text modalities (titles, captions, and image tags) and that were identified as English by the Python `langid` package. Lastly, due to large differences in the number of samples per class and the infeasibility of using all 999 classes, we selected 103 of the ImageNet-1K classes and re-labeled the images into 86 basic-level class categories based on semantic relatedness, roughly balancing the number of images per category. If a sample’s text contained its new label, we replaced the occurrence with the text [OMIT]. After imposing new class labels, we split the images for each class into training, validation, and test pools using a 50/10/40 split. In our simulations, we selected a few training examples for each class from the training pool, reserved the validation pool for hyperparameter tuning (e.g., prompt wording), and reported results on a subset of the test pool due to cost limitations.

Implementations of Cognitive Architectures

Here, we provide implementation details of how we ablated various components to instantiate the cognitive architectures found in Fig. 3.

Vision-Knowledge-Reasoning (Fig. 3A). We originally attempted to implement the ideal architecture in Fig. 1 using GPT-4V but encountered limits on the number of images per prompt using our API, making this infeasible. Instead, we removed examples and implemented the cognitive architecture in Fig. 3A by querying a VLM, GPT-4V, to classify images without providing any examples prior to testing.




Figure 3: Cognitive architectures used in simulations. **Top Left:** A full vision language model but lacking examples. **Top Right:** Possible architectures when we remove language. **Bottom:** Possible architectures when removing vision. The bottom left box shows a full LLM, and the right three represent removing one component at a time: examples, knowledge, or reasoning. Colored boxes are used to differentiate these models in subsequent presentation of the results.

Knowledge-Reasoning-Examples (Fig. 3D). We then removed visual processing from the VLM, shown by Fig. 3D, by using its language backbone: GPT-4¹. To use all components, including examples, we queried GPT-4 by providing 3 text examples per class with their corresponding labels prior to testing, using the tags associated with each image as semantic text representations of the image. The next three architectures each ablate one component from this LLM.

Knowledge-Reasoning (Fig. 3E). To ablate examples, we simply removed examples from the prompt for GPT-4. This is similar to Vision-Knowledge-Reasoning (Fig. 3A), but with tags instead of images. By removing training examples, we prohibited the model from sampling the data distribution.

Reasoning-Examples (Fig. 3F). Knowledge cannot be cleanly removed due to its distributed nature in LLMs, so instead we limited its usefulness. To do so, we provided GPT-4 with 3 examples per class before test time, similar to Knowledge-Reasoning-Examples (Fig. 3D), but we use fake (i.e. non-English) words with a 1:1 mapping to the true class labels. Because the fake word labels bear no semantic significance and break the semantic connection between the tags and labels, this setup *diminishes* the utility of prior knowledge. The pool of fake words was generated by GPT

Knowledge-Examples (Fig. 3G). Because reasoning cannot be entirely removed, we reduced it to be trivial—where our threshold for non-trivial reasoning is adapting to new tasks

without adding or modifying weights. To do so, we replaced GPT-4 with OpenAI’s text-embedding-ada-002 (Greene et al., 2015) (henceforth referred to as Ada) embedding model with a linear classifier. We used Ada to obtain embeddings of each sample’s tags, then trained a linear classifier with 3 examples per class. Because the Ada-based linear classifier must learn parameters to perform classification, while GPT does not, this model has trivial reasoning capabilities at best.

Vision-Knowledge-Examples (Fig. 3B). Now, we outline the implementations of cognitive architectures where vision is isolated. However, because OpenAI did not release the vision model of GPT-4V, we used a ResNet50 (He et al., 2016) classifier as a proxy. The model was pre-trained on the PASS image dataset (Asano et al., 2021), equipping it with knowledge. We recognize that the scale of the data and model is much smaller than that of GPT-4V, but we leave this limitation for future work. Similar to Knowledge-Examples (Fig. 3G), we used the pre-trained model to obtain embeddings, now from images, then trained a linear classifier using 3 images per class. The ResNet50 classifier lacks reasoning for the same reason as the Ada classifier.

Vision-Examples (Fig. 3C). We removed prior knowledge from the vision-only architecture by randomly initializing our ResNet50 classifier (i.e., no pre-training). We used the 3 samples per class to train the full model from scratch, effectively removing information beyond the small number of examples.

Additional Implementation Details. Prompts for models with reasoning (i.e., GPT-4V and GPT-4) contained the following: task instructions, the list of classes, shuffled training examples with our class labels (if providing examples), and a new sample. We accessed language models using the Azure

¹The model can only receive text input, so we consider it an LLM. However, GPT-4 may have been trained on both images and text. Additionally, GPT-4V technically uses GPT-4-Turbo, a smaller variant of GPT-4, as a backbone. However, preliminary experiments comparing the two showed similar performance.




Figure 4: Results from simulations. The horizontal axis lists components present in each setting where **bold** indicates component present and *gray and italicized* indicates component missing. Vision, knowledge, and reasoning are part of a model’s architecture and examples refers to the model seeing relevant examples of the new task. The highest bar in dark blue titled Vision-Knowledge-Reasoning is the full VLM. To the right we isolated language and its various components, and to the left, we isolated components of vision. Bar colors correspond to the boxes around models shown in Fig. 3. Error bars represent standard error of the mean calculated by treating each trial as a binomial random variable.

OpenAI Python API with the temperature set to 0 to encourage consistency across responses. When the prompt length was too long for GPT-4’s 8K token limit, we used GPT-4-32K with a 32K token limit. Linear classifiers were implemented using the `sklearn.neural_network.MLPClassifier` package in Python with no hidden layer, and we performed a hyperparameter search to determine the learning rate. The pre-trained ResNet50 was obtained from the Pytorch model repository. Models were evaluated on a test set of $n = 2500$ samples using the 86 basic-level class categories we defined.

Simulations

We conducted simulations to address three questions:

Simulation 1: How close can an LLM (lacking visual processing) get to the performance of a VLM (with visual processing)? We compared the full Vision-Knowledge-Reasoning (Fig. 3A) VLM to the full Knowledge-Reasoning-Examples (Fig. 3D) LLM. The goal was to obtain a lower bound on the value of language in visual understanding by removing visual processing.

Simulation 2: How do the LLM’s various components contribute to its performance? Here, we used the full Knowledge-Reasoning-Examples (Fig. 3D) LLM as a baseline and compared it to the following three cognitive architectures, each removing one of the components:

Knowledge-Reasoning (Fig. 3E), Reasoning-Examples (Fig. 3F), and Knowledge-Examples (Fig. 3G). By removing one component at a time, we could evaluate how necessary each is to understanding in the absence of visual processing.

Simulation 3: How does a vision model lacking language compare to the VLM? Lastly, we compared the vision-only architectures, Vision-Knowledge-Examples (Fig. 3B) and Vision-Examples (Fig. 3C) to the full VLM Vision-Knowledge-Reasoning (Fig. 3A). The purpose was to understand how visual processing in the absence of language could perform, helping us identify a lower bound of both the necessity of language in visual understanding and of knowledge when lacking language.

Results

Results from all simulations can be found in Fig. 4.

Simulation 1: Here, we compared a VLM without examples to a full LLM with examples. In Fig. 4, the VLM performance is the dark blue bar labeled Vision-Knowledge-Reasoning and the full LLM is the red bar in the center labeled Knowledge-Reasoning-Examples. The VLM baseline obtains an accuracy of 90.1% (Standard Error of the Mean (SEM)=0.60%) while the LLM only reaches 69.84% (SEM=0.92%, $p < 0.001$), maintaining just over 75% of the VLM’s performance.

Simulation 2: We ablated the language components of a full LLM, demonstrated by moving to the right in Fig. 4. P-values are reported compared to the full LLM from Simulation 1. As shown by the orange bar labeled *Reasoning-Examples* on the far right, when limiting access to prior knowledge by replacing true labels with fake words, the language model performance drops from 69.84% to 37.00% (SEM=0.97%, $p < 0.001$). When eliminating training examples, (see the yellow bar labeled *Knowledge-Reasoning*), model performance drops to 43.10% (SEM=0.99%, $p < 0.001$). Lastly, when reducing non-trivial reasoning from the language model, performance drops to 45.00% (SEM=0.99%, $p < 0.001$), demonstrated by the pink bar labeled *Knowledge-Examples*.

Simulation 3: We removed language entirely, demonstrated by the bars on the left side of the VLM performance. We report p-values compared to the VLM performance. When the vision-only model has prior knowledge, performance drops to 45.60% (SEM=1.00%, $p < 0.001$) as seen by the light blue bar labeled *Vision-Knowledge-Examples* in Fig. 4. However, when prior knowledge is removed, the model only achieves 0.84% (SEM=0.18%, $p < 0.001$) – approximately random chance. This is seen by the gray bar labeled *Vision-Examples* on the far left.

Discussion

Our work aims to investigate the contribution of aspects of language and visual processing to understanding the visual world. Based on our results, visual processing is a large component of visual intelligence, but sophisticated LLMs—when equipped with strong prior knowledge, reasoning mechanisms, and a few training examples—perform surprisingly well. However, only when all three are present does the model without visual processing perform remotely close to a joint vision and language system. When removing any one of these three, performance drops significantly and by similar amounts compared to removing any other component. These observations suggest that all three—knowledge, reasoning, and examples—are necessary and have approximately equal importance for visual understanding via textual inputs.

On the vision side, a vision model with knowledge but no reasoning achieves approximately half the performance of the VLM, which is comparable to a language model missing any one component. However, removing prior knowledge reduces performance to essentially random chance. This suggests that, when lacking sophisticated reasoning, sufficient prior knowledge is necessary for visual understanding.

Limitations and Future Work

While the results of our simulations provide insight into the roles of language and vision in learning from limited data, we acknowledge a few limitations. First, because the text we use as our language input was obtained from the image’s source rather than annotated, they vary in descriptiveness. Future work should explore collecting new tags via crowdsourcing processes like STEP-Tag (Marjeh, Van Rijn, et al., 2022)

or by using a VLM. We also acknowledge the limitations of solely running our evaluations on GPT due to the closed nature of its architecture and training data. To increase the robustness of our findings, we aim to repeat experiments on open-source models such as LLaMA (Touvron et al., 2023) and LLaVA (Liu et al., 2023). Lastly, as mentioned earlier, we acknowledge the pre-trained image model is smaller than the VLM and uses a much smaller dataset, which may be limiting its performance (and hence underestimating the contribution of prior knowledge in the vision case). We hope to find or train image models on comparable amounts of data to further investigate our hypotheses.

Conclusion

New AI models bring with them the opportunity to understand the factors that give rise to complex and intelligent behavior. Here, we used VLMs to analyze the contributions of visual processing and language to visual understanding and further dissected language into more granular components. Our findings show that a LLM with prior knowledge, reasoning, and a few relevant examples can surprisingly recover three-quarters of a powerful VLM’s performance, despite not having access to the actual images. Further, the synergy of all three language components mentioned above is necessary to achieve this performance; removing any one severely hinders the model’s abilities. Lastly, vision models lacking language fall behind full language models and do not significantly outperform language models missing one of the key mechanisms. This suggests that a language of thought (Fodor, 1975) may indeed be valuable even in perceptual tasks like visual classification. While these findings do not directly shed light on how humans—a very different kind of system—use vision and language, they provide a lower bound on the signal that can be derived from these sources.

In addition to identifying language’s role in visual understanding, this work also contributes to demystifying high-performing VLMs and LLMs by applying cognitive science approaches to understanding the powerful black boxes. While we are far from understanding VLMs and LLMs in their entirety, we believe that the approach we have developed here—specifying a cognitive architecture and exploring its structure through a series of ablations—is a powerful way to begin identifying the key components of these systems and ultimately increasing their transparency.

Acknowledgements. This material is based upon work supported by Microsoft Azure credits supplied to Princeton, by a Microsoft Foundation Models grant, and by the National Science Foundation under Grant No. 2107048 and 2145198. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. We would also like to thank (in alphabetical order by last name) Jihoon Chung, Amaya Dharmasiri, Max Gonzalez Saez-Diez, Ryan Liu, Maya Malaviya, Shruthi Santhanam, Xindi Wu, and William Yang for their input and feedback.

References

- Ahn, M., Brohan, A., Brown, N., Chebotar, Y., Cortes, O., David, B., ... Zeng, A. (2022). Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. *arXiv preprint arXiv:2204.01691*.
- Aleman, A., Van Lee, L., Mantione, M. H., Verkoijen, I. G., & de Haan, E. H. (2001). Visual imagery without visual experience: evidence from congenitally totally blind people. *Neuroreport*, 12(11), 2601–2604.
- Anderson, J. R. (2013). *The Architecture of Cognition*. Psychology Press.
- Asano, Y. M., Rupprecht, C., Zisserman, A., & Vedaldi, A. (2021). PASS: An ImageNet replacement for self-supervised pretraining without humans. *NeurIPS Track on Datasets and Benchmarks*.
- Bansal, T., Jha, R., Munkhdalai, T., & McCallum, A. (2020, November). Self-Supervised Meta-Learning for Few-Shot Natural Language Classification Tasks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (pp. 522–534). Association for Computational Linguistics.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... Amodei, D. (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Chopra, S., Hadsell, R., & LeCun, Y. (2005). Learning a Similarity Metric Discriminatively, with Application to Face Verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '05)* (Vol. 1, pp. 539–546).
- Connolly, A. C., Gleitman, L. R., & Thompson-Schill, S. L. (2007). Effect of congenital blindness on the semantic representation of some everyday concepts. *Proceedings of the National Academy of Sciences*, 104(20), 8241–8246.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., ... Houlsby, N. (2021). An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Driess, D., Xia, F., Sajjadi, M. S. M., Lynch, C., Chowdhery, A., Ichter, B., ... Florence, P. (2023). PaLM-E: An Embodied Multimodal Language Model. *arXiv preprint arXiv:2303.03378*.
- Fang, A., Ilharco, G., Wortsman, M., Wan, Y., Shankar, V., Dave, A., & Schmidt, L. (2022). Data Determines Distributional Robustness in Contrastive Language Image Pre-Training (CLIP). In *International Conference on Machine Learning* (pp. 6216–6234).
- Feng, Y., Lin, J., Dwivedi, S. K., Sun, Y., Patel, P., & Black, M. J. (2023). PoseGPT: Chatting about 3D Human Pose. *arXiv preprint arXiv:2311.18836*.
- Fodor, J. A. (1975). *The Language of Thought* (Vol. 5). Harvard University Press.
- Gao, T., Fisch, A., & Chen, D. (2021, August). Making Pre-trained Language Models Better Few-shot Learners. 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)* (pp. 3816–3830). Association for Computational Linguistics.
- Greene, R., Sanders, T., Weng, L., & Neelakantan, A. (2015). *New and Improved Embedding Model — OpenAI.com*. <https://openai.com/blog/new-and-improved-embedding-model>. ([Accessed 24-01-2024])
- Guo, Z., Dong, B., Ji, Z., Bai, J., Guo, Y., & Zuo, W. (2023). Texts as Images in Prompt Tuning for Multi-Label Image Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 2808–2817).
- Hadsell, R., Chopra, S., & LeCun, Y. (2006). Dimensionality Reduction by Learning an Invariant Mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR '06)* (Vol. 2, pp. 1735–1742).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778).
- Heinzerling, B., & Inui, K. (2021, April). Language Models as Knowledge Bases: On Entity Representations, Storage Capacity, and Paraphrased Queries. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume* (pp. 1772–1791). Association for Computational Linguistics.
- Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., ... Duerig, T. (2021). Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision. In *International Conference on Machine Learning* (pp. 4904–4916).
- Kerr, N. H. (1983). The role of vision in “visual imagery” experiments: Evidence from the congenitally blind. *Journal of Experimental Psychology: General*, 112(2), 265–277.
- Kim, J. S., Aheimer, B., Montané Manrara, V., & Bedny, M. (2021). Shared understanding of color among sighted and blind adults. *Proceedings of the National Academy of Sciences*, 118(33), e2020192118.
- Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., ... Yan, J. (2022). Supervision Exists Everywhere: A Data Efficient Contrastive Language-Image Pre-training Paradigm. In *International Conference on Learning Representations*.
- Liu, H., Li, C., Wu, Q., & Lee, Y. J. (2023). Visual Instruction Tuning. In *NeurIPS*.
- Marjeh, R., Sucholutsky, I., Summers, T., Jacoby, N., & Griffiths, T. (2022). Predicting Human Similarity Judgments Using Large Language Models. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 44).

- Marjeh, R., Sucholutsky, I., van Rijn, P., Jacoby, N., & Griffiths, T. L. (2023). What Language Reveals about Perception: Distilling Psychophysical Knowledge from Large Language Models. In *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.
- Marjeh, R., Van Rijn, P., Sucholutsky, I., Sumers, T., Lee, H., Griffiths, T. L., & Jacoby, N. (2022). Words are all you need? Language as an approximation for human similarity judgments. In *The Eleventh International Conference on Learning Representations*.
- Newell, A., Simon, H. A., et al. (1972). *Human Problem Solving* (Vol. 104) (No. 9). Prentice-hall Englewood Cliffs, NJ.
- OpenAI. (2023). GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.
- Petroni, F., Rocktäschel, T., Riedel, S., Lewis, P., Bakhtin, A., Wu, Y., & Miller, A. (2019, November). Language Models as Knowledge Bases? 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)* (pp. 2463–2473). Hong Kong, China: Association for Computational Linguistics.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., ... Sutskever, I. (2021, 18–24 Jul). Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning* (Vol. 139, pp. 8748–8763). PMLR.
- Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al. (2018). Improving Language Understanding by Generative Pre-Training.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language Models are Unsupervised Multitask Learners. *OpenAI blog*, 1(8), 9.
- Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., ... Lu, J. (2022). DenseCLIP: Language-Guided Dense Prediction with Context-Aware Prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 18082–18091).
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., ... Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3), 211–252. doi: 10.1007/s11263-015-0816-y
- Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., & Kiela, D. (2022). FLAVA: A Foundational Language And Vision Alignment Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 15638–15650).
- Sumers, T. R., Yao, S., Narasimhan, K., & Griffiths, T. L. (2023). Cognitive Architectures for Language Agents. *arXiv preprint arXiv:2309.02427*.
- Sun, R. (2024). Can A Cognitive Architecture Fundamentally Enhance LLMs? Or Vice Versa? *arXiv preprint arXiv:2401.10444*.
- Sun, X., Hu, P., & Saenko, K. (2022). DualCoOP: Fast Adaptation to Multi-Label Recognition with Limited Annotations. *Advances in Neural Information Processing Systems*, 35, 30569–30582.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., ... Lample, G. (2023). LLaMa: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is All You Need. *Advances in Neural Information Processing Systems*, 30, 5998–6008.
- Wu, N., Kera, H., & Kawamoto, K. (2023). Zero-Shot Action Recognition with ChatGPT-Based Instruction. In *International Workshop on Advanced Computational Intelligence and Intelligent Informatics* (pp. 18–28).
- Xu, Q., Peng, Y., Wu, M., Xiao, F., Chodorow, M., & Li, P. (2023). Does conceptual representation require embodiment? insights from large language models. *arXiv preprint arXiv:2305.19103*.
- Zhao, Q., Zhang, C., Wang, S., Fu, C., Agarwal, N., Lee, K., & Sun, C. (2023). AntGPT: Can Large Language Models Help Long-term Action Anticipation from Videos? *arXiv preprint arXiv:2307.16368*.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., ... Stoica, I. (2023). Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Zimler, J., & Keenan, J. M. (1983). Imagery in the congenitally blind: How visual are visual images? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(2), 269–282.