

Does Vision Accelerate Hierarchical Generalization in Neural Language Learners?

Tatsuki Kuribayashi and Timothy Baldwin

MBZUAI

tatsuki.kuribayashi@mbzuai.ac.ae

Abstract

Neural language models (LMs) are arguably less data-efficient than humans from a language acquisition perspective. One fundamental question is why this human–LM gap arises. This study explores the advantage of *grounded* language acquisition, specifically the impact of visual information — which humans can usually rely on but LMs largely do not have access to during language acquisition — on syntactic generalization in LMs. Our experiments, following the poverty of stimulus paradigm under two scenarios (using artificial vs. naturalistic images), demonstrate that if the alignments between the linguistic and visual components are clear in the input, access to vision data does help with the syntactic generalization of LMs, but if not, visual input does not help. This highlights the need for additional biases or signals, such as mutual gaze, to enhance cross-modal alignment and enable efficient syntactic generalization in multimodal LMs.

1 Introduction

Neural language models (LMs) have accelerated progress in natural language processing (NLP), but there remains a significant disparity in their data efficiency compared to humans. For instance, GPT-3 (Brown et al., 2020) is trained on approximately 2,000 times more text than a 10-year-old child is exposed to (Warstadt and Bowman, 2022) and this gap is even greater in modern large LMs, and yet the model still struggles with some language tasks. We investigate what kind of differences between human and LM language acquisition scenarios can potentially close the gap in data efficiency, specifically to achieve syntactic generalization.

One general criticism of neural LMs is their lack of grounding (Roy and Reiter, 2005; Barsalou, 2008): they learn language solely based on text and do not model the explicit association between linguistic expressions and the associated

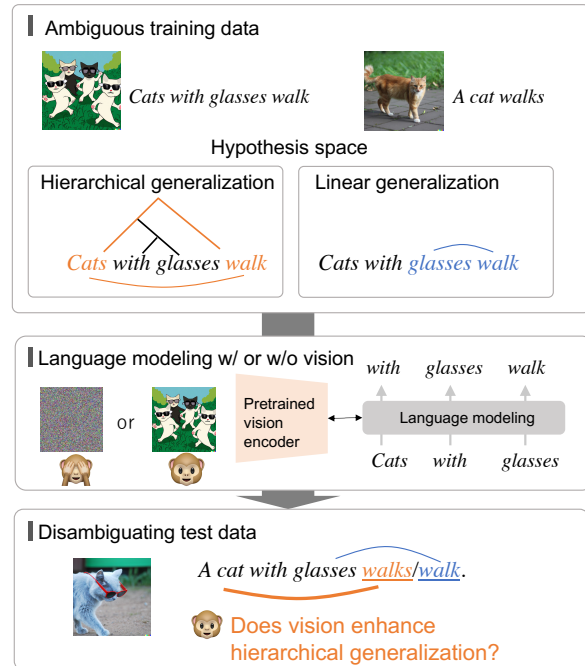


Figure 1: Overview of the experimental design. A vision-language neural model is trained on ambiguous data for a particular linguistic rule. Then, we test whether the model learned a cognitively plausible rule using data disambiguating the model’s generalization. Through this experimental scheme, we adjust whether/how the visual information helps the model infer the proper linguistic generalization.

objects/events in the real world. This naturally leads to the hypothesis that the human–LM data efficiency gap comes from this disconnect.

In this study, we investigate whether visual information, as a representative modality promoting grounding, can accelerate the emergence of the syntactic hierarchical generalization ability of LMs, which underlies human language acquisition (Chomsky, 1964). Our experiments extend the single modality version of the *poverty of stimulus* (POS) setting (Wilson, 2006; Perfors et al., 2011; McCoy et al., 2018, 2020; Warstadt and Bowman, 2020; Yedetore et al., 2023) into the vision-

and-language domain. That is, we train LMs on *ambiguous* image–text pairs in terms of particular linguistic rules (e.g., HIERARCHICAL vs. LINEAR English subject–verb number agreement rules; see Figure 1). Then, we investigate whether visual input efficiently guides the models to make cognitively plausible (hierarchical) generalizations given ambiguous data, compared to text-only models.

To adjust the visual conditions, we base our experiments on either (i) realistic image–caption data (Sharma et al., 2018), or (ii) simplified, artificial data, which is a proxy for externally-guided attentional focus. Notably, it has been argued that either strong inductive bias or additional signals, such as mutual gaze, pointing, or other forms of attentional focus, are needed to make use of multimodal input for linguistic generalization (Qu and Chai, 2008; Johnson et al., 2012) since merely adding an input modality may incur many superficial correlations and complicate rather than simplify the task (Gleitman and Gleitman, 1992; Dupoux, 2018). Thus, our investigation using the two types of multimodal data can be seen as an evaluation of the inductive bias of neural LMs toward multimodal linguistic generalization with and without such additional signals. Most work on grounded and situated multimodal LM as well as human language acquisition has focused on word learning (Hill and Wagovich, 2020; Ma et al., 2023). In this work, we extend these investigations to the acquisition of syntactic hierarchical generalizations, the central topic toward the POS setting in NLP (McCoy et al., 2018, 2020), with multimodal LMs.

In a realistic setting, we found that overall: (i) vision data **does not substantially** accelerate hierarchical generalization; (ii) this trend is consistent among 20 model settings; and (iii) this is also consistent across four different degrees of ambiguity. In contrast, with simplified, artificial data, where visual/linguistic concepts are already abstracted and simplified, we generally found the opposite trend: vision data **did** boost hierarchical linguistic generalization. These contrasts suggest that neural models have the potential to make use of visual input for linguistic generalization when the visual input is made salient either through inductive bias or external signals. However, efficient generalization via more complex and ambiguous visual input is not possible in the model variants tested either because the visual processing module lacks appropriate inductive bias or the external signals of attentional salience are absent.

2 Background

2.1 Inductive bias in language acquisition

In general, a unique generalization or rule cannot be determined solely based on the observation of finite data. The choice depends on the inductive biases of the model, such as a learner’s prior knowledge (Mitchell, 1980).

In humans: In the context of language acquisition, it has long been argued that human learners possess a strong inductive bias due to rapid language acquisition from limited language exposure (Chomsky, 1980; McCoy et al., 2018). The main question is what type of biases humans have and where these biases originate. Regarding the former question, it has been reported that humans have a bias to prefer hierarchical generalization over linear generalization in situations like those depicted in Figure 1 (Crain and Nakayama, 1987; Legate and Yang, 2002). As for the latter question, there are two primary potential sources of inductive biases: innate factors and environmental/empirical factors. To address this question, this study investigates the influence of a specific environmental factor — access to visual information during language acquisition — through computer simulations.

In neural models: Neural models typically exhibit non-human-like generalizations, such as the use of superficial cues and linear rules, as widely observed across various NLP domains (McCoy et al., 2019; Warstadt and Bowman, 2020; Warstadt et al., 2020b; McCoy et al., 2020). Large amounts of data are required to overcome such cognitively implausible biases during training (Warstadt and Bowman, 2020; Warstadt et al., 2020b). In this context, addressing the inadequate biases in neural models and tackling their data-inefficiency issues are two aspects of the same problem. Our interest lies in understanding whether and how visual information contributes to the development of appropriate inductive bias in neural language learners.

2.2 Hypotheses on the advantage of vision

There has already been some investigation into the contribution of vision in language learning. It is important to note that this study does not take a strong position on the benefits of vision but rather conducts an exploratory investigation.

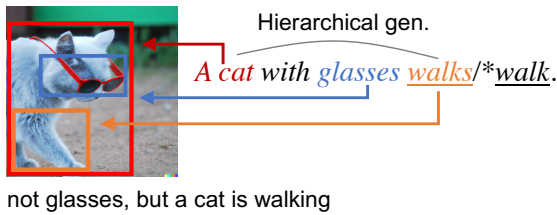


Figure 2: Images can explicate the subject–verb dependency. If a learner can ground *cat*, *glasses*, and *walk* to their visual components, they can disambiguate that what is walking is not *glasses* but *cat*; such information will potentially bias the learner’s language acquisition in favor of the linguistically correct rule.

Positive view: The general advantages of input beyond text modality in language acquisition have been historically emphasized (Goldberg, 2005; Bender and Koller, 2020). From an NLP perspective, the advantage of visual information typically for syntactic parsing was demonstrated (Shi et al., 2019; Kojima et al., 2020). Note that such NLP research used a specially-designed parser that already has a strong inductive bias (e.g., the training objective is parsing); our question is whether even vanilla neural models, a domain-general learner, with next-word prediction can take advantage of visual information for syntactic hierarchical generalization. Moreover, in achieving hierarchical generalizations in settings like that illustrated in Figure 1, intuitively, images have the potential to boost correct generalization. For example, in a sentence such as *a cat with glasses walks*, the information that it is the *cat*, not the *glasses* that is walking, could potentially bias the learning towards a hierarchical generalization. Such a clue — it is the *cat* walking and not the *glasses* — would be explicit in the image (Figure 2) if the learner or model understands the visual concepts of *cat*, *glasses*, *walk*, and their composition (e.g., *walking cat*). In addition, at least for the number agreement problem, the number information is, more or less, salient in the vision domain. When the number of visual objects corresponding to grammatical subjects changes, the content of the image will change drastically, while in the text domain, only a few characters/tokens are changed.¹

Negative view: There is also skepticism that merely providing visual input without appropri-

¹Strictly speaking, grammatical and physical (visual) numbers are not exactly the same concepts (Spector, 2007; Zweig, 2009).

ate linguistic knowledge or attentional focus could over-complicate the problem, e.g., increase the potential for superficial correlations (Gleitman and Gleitman, 1992; Dupoux, 2018). For example, Gleitman and Gleitman (1992) and McDonough et al. (2011) assumed that children use syntactic category information to ground words to visual input; this implies that syntactic knowledge comes first, followed by grounding. These studies generally claim that the advantage of input beyond text in language acquisition could be driven by both humans’ prior knowledge and visual input. In this sense, if neural LMs, which are assumed to have no innate knowledge, fail to accelerate linguistic generalization with visual input, this implicitly highlights the necessity of specific learners’ inductive biases or additional attentional signals in multimodal language acquisition. Beyond syntactic generalization, there are actually some reports that visual input does not enhance the fundamental linguistic knowledge of models (Yun et al., 2021; Wang et al., 2023) or classifiers (Ma et al., 2021) (c.f. contemporaneous work by Zhuang et al. (2024) arguing multimodal input does accelerate neural LM word learning on some smaller datasets).

Similar attempts: Concurrent works have empirically investigated what linguistic ability particular neural networks can acquire solely from developmentally-plausible multimodal data that is recorded by a head-mounted camera of English-speaking children (Vong et al., 2024; Qin et al., 2024; Wang et al., 2023), motivated by the general, historical debates on the empiricism toward language acquisition (Elman, 1990; Kirov and Cotterell, 2018). Although their results suggest the learnability of certain linguistic properties by image-captioning models and these data, the exact advantage of visual input itself was nuanced on BLiMP (Wang et al., 2023), beyond the focus (Qin et al., 2024), or unclear (Vong et al., 2024) since the evaluation tasks are image-classification/mapping, where it is somewhat obvious to see the advantage of visual input. Furthermore, these studies examined a very limited variant of visual encoders; thus, the generality of the results was unclear. Our evaluation potentially achieves fairer comparisons since the task itself (acceptability judgment toward syntactic generalization) is agnostic to the existence of visual modality, and we observe generally consistent results from 12 variants of vision-language models.

3 Problem definition

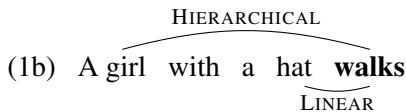
We briefly introduce the poverty of stimulus (POS) settings (Wilson, 2006; Perfors et al., 2011; McCoy et al., 2018, 2020; Warstadt et al., 2020b; Warstadt and Bowman, 2020, 2022; Yedetore et al., 2023). Through our experiments, we aim to quantify whether vision accelerates cognitively-plausible generalization in neural LMs.

3.1 HIERARCHICAL vs. LINEAR generalizations

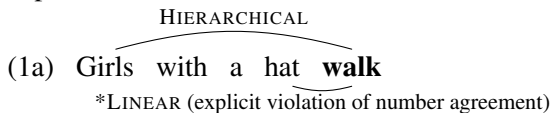
We use the subject–verb number agreement rule as a target phenomenon. In English, the subject and corresponding verb should match in terms of their grammatical number:

- (1) a. **Girls** with a hat **walk**.
- b. A **girl** with a hat **walks**.

Here, Example (1b) is *ambiguous* because a learner can infer at least two different generalizations from this example alone, i.e., HIERARCHICAL and LINEAR rules:



The HIERARCHICAL rule associates the grammatical number of a verb with that of its grammatical subject, while the linear one associates the number between a verb and its closest noun in a linear word order. By contrast, Example (1a) is not ambiguous in terms of the HIERARCHICAL and LINEAR rules since the number does not match under the LINEAR assumption:



Our interest lies in which rule a particular learner acquires from *ambiguous* data and what factors (e.g., vision) can guide the learner to prefer the HIERARCHICAL rule that is linguistically correct (Section 3.2). The motivation for this experimental setting is further described in Section 3.2.

We only employed this subject–verb number agreement setting in our experiments, although other studies have focused on different syntactic transformation tasks, such as question formulation or passivization (McCoy et al., 2020; Warstadt and

Bowman, 2020; Mueller et al., 2022). Our motivation is the ease of collecting natural images for sentences with subject–verb agreement and the strong correlations between image entities and grammatical number. Such correlations are either absent or weak in the case of interrogative vs. declarative sentences and passive vs. active mood.

3.2 Poverty of stimulus setting

Children acquire HIERARCHICAL rules despite the scarcity of disambiguating sentences, like Example (1a), in real language exposure (Crain and Nakayama, 1987; Legate and Yang, 2002). Building on this scenario, we expose a model to (nearly) ambiguous data where the generalization cannot be determined as to whether LINEAR or HIERARCHICAL rules are correct. Then, we evaluate the model in terms of which rule is obtained from the ambiguous data via a test using unambiguous data.

Data splitting strategy: We split data into two groups: (i) those that do not disambiguate LINEAR and HIERARCHICAL rules (AMBIGUOUS); and (ii) those that support the HIERARCHICAL rule (UNAMBIGUOUS). Examples are shown in Table 1. Basically, the AMBIGUOUS instances are used in training, and UNAMBIGUOUS instances are used in evaluation. We insert a few held-out UNAMBIGUOUS instances into training data since it is counter-intuitive that a learner *never* encounters direct evidence for hierarchical generalizations, i.e., UNAMBIGUOUS instances, during language acquisition. Therefore, we controlled the injection rate — the extent to which disambiguating data appear during training — for experiments analyzing sensitivity to the scarcity of direct evidence (Section 4.1).

Model comparison: In this series of experiments, we compare neural models that can access visual information (👁️) and ones that do not (👂) to assess the contribution of vision modality. Note that “visual information” in this study denotes an image compatible with the meaning of a sentence, i.e., we use image–caption pairs. The source of image caption data is described in Section 3.3.

3.3 Data

We introduce two complementary data types: (i) NATURAL captions; and (ii) ARTIFICIAL captions. The NATURAL captions are collected from an image–caption corpus, while the ARTIFICIAL captions are automatically created by rules to simplify the task.







Setting Split		NATURAL		ARTIFICIAL
AMBIGUOUS		<i>girl aged stands with a hand on a tree alone</i>		<i>a lime rectangle with a red rectangle waves its hand</i>
		<i>young boys with school uniforms and backpacks prepare for school on an early morning</i>		<i>two yellow circles with three blue hexagons take a photo</i>
DISAMBIGUATING		<i>young girls dressed in colonial gear tie their shoes at farm</i>		<i>two red rectangles with a black circle play soccer</i>

Table 1: Examples of image-caption pairs. The NATURAL data is collected from conceptual captions corpus, and the ARTIFICIAL data is generated by rules. In the AMBIGUOUS set, the grammatical numbers of **verb**, its corresponding **subject**, and its immediately preceding **noun** are identical; in this sense, they are ambiguous toward which is the correct rule of number agreement, LINEAR or HIERARCHICAL. By contrast, the DISAMBIGUATING instances disambiguate the rule.

NATURAL dataset: We extracted image-caption pairs from the Conceptual Captions Corpus (Sharma et al., 2018), which is a widely-used and relatively large-scale image-caption dataset. Specifically, we first collected captions that: (i) form a complete sentence, (ii) do not have grammatical errors²; and (iii) do not have collective expressions such as *family* or *pair of* since these are confusing in terms of grammatical number. Then, we split the data into the AMBIGUOUS and UNAMBIGUOUS sets using a dependency parser.³ Note that there might be parsing errors in this process, but we later observe that the models did not prefer the HIERARCHICAL rule without injection of any disambiguating examples; this suggests that such errors do not inadvertently bias the model toward the HIERARCHICAL rule. Examples are shown in the left part of Table 1. The training set (AMBIGUOUS part) consists of 348,861 image-caption pairs, and the unambiguous test set consists of 1,253 pairs.

ARTIFICIAL dataset: Image-caption pairs were generated by rules. Specifically, a caption is first generated with the template of NUM1 COLOR1 SHAPE1 with NUM2 COLOR2 SHAPE2 VP; then, the corresponding image is automatically created (the detailed process is shown in Appendix A). Examples are shown in the right part of Table 1. As with the NATURAL setting, we split the data into AMBIGUOUS and UNAMBIGUOUS cases. Then, training and test data are created with different in-

jection rates. The training set (AMBIGUOUS part) consists of 15,000 pairs, and the test set consists of 5,000 pairs.

This setting limits the variations of linguistic/visual concepts and sentence constructions compared to the NATURAL setting, and importantly, the alignment between linguistic and visual components can easily be extracted since the image only has visual objects related to the caption (less confounding factors), and word types and visual features have a one-to-one relationship (no lexical ambiguity; see appendix A). Thus, we use this artificial data setting to approximate the richer environment in which learners exploit visual inductive bias, gaze recognition, pointing and other extralinguistic signals of salience and focus to interpret otherwise ambiguous linguistic input.

3.4 Evaluation

For each UNAMBIGUOUS instance, we prepared two candidate captions differing only in the verb’s grammatical number (e.g., *two red rectangles with a black circle **play/plays** soccer*); one aligns with the HIERARCHICAL rule, and the counterfactual one with the LINEAR rule by modifying the grammatical number of its main verb. The model’s generalization preference is determined by which caption has a higher probability.

Specifically, a model θ computes the probabilities of each caption $\mathbf{s} = [w_1, \dots, w_n]$ given the corresponding image v :

$$p(\mathbf{s}|v) = \prod_{t=1}^n p_{\theta}(w_t | \mathbf{w}_{<t}, v) , \quad (1)$$

²We used language-tool-python 2.7.1

³We used SpaCy (Honnibal et al., 2020).

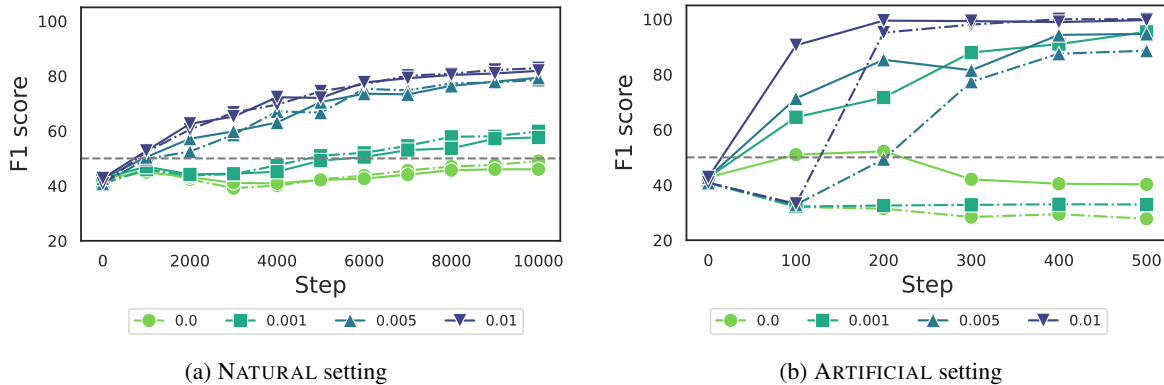


Figure 3: Generalization performance of the model initialized with Vit-base. The x -axis denotes the parameter update steps, and the y -axis denotes the preference for the HIERARCHICAL generalization rule (F1 scores multiplied by 100). We adopted four settings with different injection rates of $\{0, 0.001, 0.005, 0.01\}$. The normal lines correspond to the model with visual input (👁️), and the dashed lines correspond to the preference of those without visual input (👂). The chance rate of the F1 score is 50.

where $w_{<t}$ denotes the left context of w_t in the caption s . We calculated the macro-F1 score, considering the inflection corresponding to the HIERARCHICAL rule as correct and treating the task as a binary classification problem for selecting a grammatically-correct sentence. As we are interested in language acquisition efficiency, we report F1 scores at various intermediate training steps.

3.5 Models

We use the Transformer seq2seq image-caption model as a vision-and-language model \checkmark , with the encoder set as a pre-trained vision encoder like ViT (Dosovitskiy et al., 2021). An image is input to the encoder, and the decoder predicts the caption in a left-to-right manner, accessing visual information via cross-attention. Intuitively, this can be viewed as a sentence-level LM that can access visual information. For the image-less 👂 model, we replaced the input image with a white noise image during training and inference. Models are trained with cross-entropy loss to generate the reference caption. The vision encoder is further updated during the training.

We adopted the GPT-2 small (124M) architecture (Radford et al., 2019) for the decoder, with parameters randomly initialized, considering a language acquisition scenario from scratch. As an encoder, we initially used Vit-base (Dosovitskiy et al., 2021) in Section 4.1 and further examined various encoders in Section 4.2 to enhance the generality of the conclusion. Hyperparameters are listed in Appendix B. In each setting, we train two models with different seeds and report the average score.

4 Experiments

4.1 Generalization preferences

We first analyze the model using the pre-trained Vit-base encoder. We examined four different injection rates of $\{0, 0.001, 0.005, 0.01\}$; for example, the rate 0.001 means that ten held-out UNAMBIGUOUS instances are added into the training data if the original training data size is 10,000.

Results: The results are shown in Figure 3, with scores averaged across models with different seeds. These indicate the following:

- In the NATURAL setting, visual inputs do not generate a substantial difference in generalization efficiency.
- In the ARTIFICIAL setting, visual inputs accelerate hierarchical generalization, especially at the early stages of learning.
- At the initial stage of learning in the NATURAL and ARTIFICIAL settings with a low injection rate, the LINEAR rule emerged (F1-score below chance rate), indicating that the model originally has a LINEAR bias. This is consistent with existing studies in the text-only domain (McCoy et al., 2020).
- With moderate rates of injection, e.g., above the rate of 0.005, the models gradually acquired the HIERARCHICAL rule, showing sensitivity to the slight bias in data distribution.

We further discuss the implications of the contrasting results between the NATURAL and ARTIFICIAL settings in Section 5.

Models	Vision	NATURAL			ARTIFICIAL	
		1,000	5,000	10,000	100	500
Vit-base (86M)	✓	52.8	72.0	81.9	90.6	99.7
	Δ	+0.41	-2.38	-0.94	+57.4	-0.31
Vit-large (307M)	✓	52.9	74.9	83.1	52.6	92.2
	Δ	+0.93	-1.13	+0.65	+19.4	-7.76
Vit-huge (632M)	✓	52.6	73.9	82.6	42.6	100
	Δ	+1.98	-2.07	+0.10	+9.21	0.00
Beit-base (86M)	✓	46.7	59.0	66.4	45.8	74.8
	Δ	+2.99	+5.68	-1.50	+11.7	-25.0
Beit-large (307M)	✓	45.6	65.3	73.3	38.3	57.7
	Δ	+1.57	+4.32	+3.80	+5.09	-38.4
Deit-base (86M)	✓	54.9	72.5	81.2	67.4	99.9
	Δ	+4.23	-1.77	-1.35	+32.9	+0.08
Deit-small (22M)	✓	52.9	73.7	83.2	73.1	94.1
	Δ	+3.79	-0.16	-0.52	+27.1	-5.86
Deit-tiny (5M)	✓	52.6	73.5	81.0	88.8	87.8
	Δ	+2.16	-1.29	-1.87	+32.5	-12.2
Swin-base (88M)	✓	53.0	73.0	81.8	80.5	100
	Δ	+0.92	-2.61	-1.05	+33.2	0.00
Swin-large (197M)	✓	53.3	73.9	82.4	74.9	100
	Δ	+0.85	-0.79	-0.11	+39.3	0.00
Scratch (86M)	✓	49.3	72.6	81.0	50.7	100
	Δ	+1.75	-3.22	-1.62	+5.10	0.00
Vit-GPT2 (86M)	✓	95.6	97.0	96.6	90.8	100
	Δ	+0.04	+0.18	-0.11	-9.21	0.00

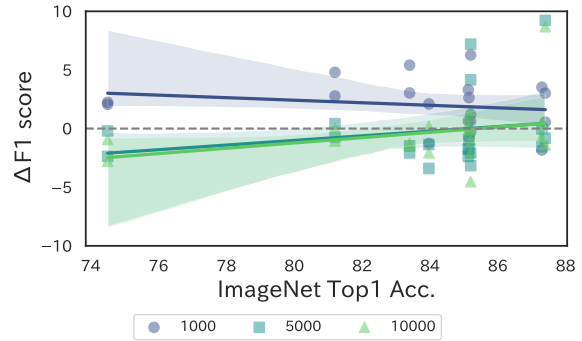
Table 2: The preference for HIERARCHICAL generalization (F1 score) of models without various configurations. F1 scores are multiplied by 100. The column names such as 1,000, 5,000, and 10,000 denote the training steps. Scores in the ✓ row indicate the results of models with visual inputs 🖼️, and those in Δ indicate the score difference between models with and without visual inputs (🖼️ - 🗣️).

4.2 Vision encoder variations

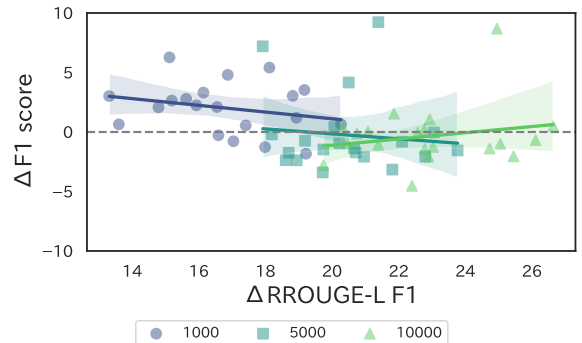
To investigate whether our results are specific to a particular model setting, we further analyze ten vision-language models with different encoder-decoder settings, demonstrating general consistency across various settings.

Generality of the (in)effectiveness of vision:

We tested the models using ten different vision encoders: Vit- $\{\text{base, large, xlarge}\}$ (Dosovitskiy et al., 2021), Beit- $\{\text{base, large}\}$ (Bao et al., 2022), Deit- $\{\text{base, small, tiny}\}$ (Touvron et al., 2021), and Swin- $\{\text{base, large}\}$ (Liu et al., 2021). We also examined two baselines: one using randomly initialized Vit-base (Scratch) and a model using the pre-trained GPT-2 (Radford et al., 2019) as a decoder (Vit-GPT2). Note that the Vit-GPT2 model



(a) Relationship between encoders’ ImageNet accuracy (x-axis) and their advantage in HIERARCHICAL generalization (F1 score difference of 🖼️ - 🗣️; y-axis). The F1 score is measured at several checkpoints during training (1000, 5000, and 10000).



(b) Relationship between encoders’ captioning performance in the validation set (x-axis) and their advantage in HIERARCHICAL generalization (F1 score difference of 🖼️ - 🗣️; y-axis). These scores are measured at several checkpoints during training (1000, 5000, and 10000).

Figure 4: Relationship between CV-oriented metrics and the contribution to HIERARCHICAL generalization in the NATURAL setting. Each dot corresponds to each setting $\{10 \text{ encoders}\} \times \{2 \text{ seeds}\} \times \{3 \text{ training steps}\}$, and its color/shape corresponds to training steps.

is already trained on large-scale text data, including disambiguating sentences; thus, it is not surprising that they achieve hierarchical generalization. We fix the inoculation rate to 0.01 in this Section.

The results are summarized in Table 2. The observations are similar to those in Section 4.1: (i) the effect size of the visual input factor is larger in the ARTIFICIAL setting than the NATURAL setting, especially at the early stage of learning;⁴ (ii) vision data generally has a positive/negative effect on the generalization at the early/late stage.⁵ Note that

⁴With a two-sided Wilcoxon rank-sum test, the Δ scores from the 100-step ARTIFICIAL setting was significantly larger than those in the 1000-step setting across models and seeds ($p = 5.3e-4 < 0.05$).

⁵With a two-sided one-sample t-test, the Δ scores were significantly larger than zero across models and seeds in the 1,000-step NATURAL setting ($p = 4.1e-4 < 0.05$) and 100-

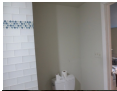

	<i>the walls over the toilet need a small cabinet</i>
	<i>boys with eyes like that drive me crazy</i>

Table 3: Examples exhibiting some challenging features of NATURAL image captions.

models with visual input (👁️) achieved ROUGE-L F1 scores of 30–40 in the NATURAL setting (Appendix B), whereas those without visual input (👂) yielded the scores of around 15; this improvement indicates that the models do not ignore visual input.

As minor points, Beit-based models yielded somewhat idiosyncratic trends (HIERARCHICAL generalization is hurt at the late stage in the ARTIFICIAL setting). In addition, as a sanity check, Vit-GPT2, which is pre-trained over a massive amount of text data, achieved almost perfect hierarchical generalization from the early stages of training in both NATURAL and ARTIFICIAL settings.

Which vision encoder relatively accelerates hierarchical generalization? Different vision encoders generally show a similar trend, but the degree of their advantage is slightly different—what kind of encoder benefits most from vision inputs? This can be viewed as an evaluation of vision encoders from a cognitive perspective. Figure 4 shows the following: (i) no clear relationship between the encoders’ ImageNet top-1 accuracy⁶ and their contribution to linguistic HIERARCHICAL generalization (Δ F1 score in Table 2); and (ii) no clear relationship between image-captioning performance and the contribution to hierarchical generalization. Note that the Δ ROUGE in Figure 4b indicates the ROUGE gain from a model without visual input to the one with visual input based on the same architecture. The results indicate that an engineeringly better vision encoder does not always lead to better linguistic generalization when combined with a language decoder.

5 Discussion and limitations

Mixed results in NATURAL and ARTIFICIAL settings: The limited advantage of vision in the NATURAL setting ($p = 1.8e-5 < 0.05$), not significant in the 5,000/10,000-step NATURAL settings ($p = 0.6$, $p = 0.4$), and lower than zero in the 500-step ARTIFICIAL setting ($p = 8.0e-3 < 0.05$).

⁶We used the scores reported in their original papers.

URAL setting suggests at least two possibilities: (i) vision is not helpful for efficient language acquisition; or (ii) vision is potentially helpful in human language acquisition scenario, but neural models lack certain human-like biases, such as learners’ prior knowledge or training/data scenario related to vision-language grounding. If one accepts the general argument about the advantage of vision and/or the advantage in the ARTIFICIAL setting as a support for the potential usefulness of visual input, vision is useful in linguistic generalization—and interpretation (ii) is plausible. Thus, the challenge lies in how the learner can extract meaningful intake from raw images and texts, and at least the modern neural models we examined might not possess such an ability. This view aligns with the considerations put forth by, for example, Gleitman and Gleitman (1992) and Dupoux (2018).

Words beyond the image content: What specific difficulties exist in the NATURAL data? One potential challenge we considered based on the dataset is that the natural caption contains information that is **not** present in the image, which might cause confusion in terms of the visual grounding of the sentence. For example, the first image in Table 3 has a caption *the walls over the toilet need a small cabinet*. In this case, the *cabinet* is not in the image, although it is not directly relevant to the subject–verb agreement. The second example’s caption in Table 3 also mentions objects beyond the image; here, the word *boys* does not refer to the boy in this image but any boy with similar eyes to him. This is potentially confusing in terms of number agreement since the grammatical subject is in plural form, but the image shows one object. These assert that visual grounding already needs linguistic knowledge and the question of where such linguistic knowledge should come from.

Coverage of the experiments: We only focused on a specific syntactic phenomenon, subject–verb number agreement rule. Extending the experimental settings to cover broader linguistic phenomena, e.g., including revisiting vocabulary acquisition (Räsänen and Khorrami, 2019), is needed to draw more general conclusions. In Appendix C, we conducted a preliminary examination using the BLiMP benchmark (Warstadt et al., 2020a) on the linguistic knowledge of models with/without vision; this also implied that visual input alone does not lead to a substantial advantage. Nevertheless, typical resources for linguistic probes, including

BLiMP, use only text input; it is not obvious how to use such data to evaluate multimodal models. We hope that this study encourages the community to build a dataset to probe the fine-grained linguistic knowledge of multimodal models.

6 Conclusions

We conducted two complementary experiments — a noisy, realistic setting and a simplified, artificial one — to investigate the advantage of vision in the syntactic generalization of LMs. Our results showed that vision accelerates proper linguistic generalization under a simplified setting, but LMs struggled with proper generalization based on noisy, realistic data. These mixed results suggest several possibilities; for example, an image can potentially boost language acquisition, but neural learners may require additional visual/linguistic prior knowledge or externally-provided attentional focus to robustly make use of *raw* images for efficient language acquisition.

Limitations

In addition to the limitations of our work raised in § 5, the following are potential concerns. First, the data size is relatively small; the training data in the NATURAL setting consists of around 3.5M tokens. Nevertheless, experiments with similar motivations have been conducted with the same or smaller scale of dataset (Nikolaus et al., 2019; Wang et al., 2023). Furthermore, at least based on the report that human infants around 18 months learn syntactic dependencies (Perkins and Lidz, 2021) and they are typically exposed to 2–7M words per year (Gilkerson et al., 2017), our data size may not be too small to learn syntactic rules.

Second, we only focused on a specific type of vision-language model—image-captioning models. There are other formulations involving vision-and-language interaction, such as text-to-image models (Ramesh et al., 2021), discrimination models like CLIP (Radford et al., 2021), or more generally, LMs with a visual input support (Alayrac et al., 2022; OpenAI, 2023). Investigating the inductive bias related to such architectural/task differences would be an interesting direction for future work. Evaluating larger models will also provide us with insights into scaling laws in this context. Having said that, such experiments require more computing resources than a typical laboratory has, which was an unrealistic direction for us to explore. More

generally, humans see both static and dynamic input during language acquisition. Therefore, extension from image to video is an important future direction of research.

Third, there are concurrent endeavors to examine the contribution of visual information to proper linguistic generalizations of neural LMs from cognitively-motivated perspectives (Wang et al., 2023; Zhuang et al., 2024); the closest initiative would be the 2nd-round of the BabyLM shared task, which includes multimodal data (Choshen et al., 2024). Enhancing the connection to such recent works will be the target of future work, and we would like to highlight that our study has employed a control to the training data properties to gain rich insights into the model’s inductive biases, which has rarely been achieved in existing multimodal experiments and is orthogonal to the holistic evaluation of pretrained vision-language models.

Ethical concerns

This study employed a widely-used, publicly available image–caption dataset, to avoid ethical concerns. In our argument, we assumed that humans usually have access to visual information during language acquisition; this is not intended to discriminate against vision-impaired people. Our general interest is in grounding, which can also be established by other modalities, and we focus on the vision modality as one case study. Perhaps our results of no advantage of visual input may be supported by the success of human language acquisition regardless of their congenital blindness; such a broader connection to human language acquisition should be enhanced in future work.

Acknowledgement

This work was partially supported by JST CREST Grant Number JPMJCR20D2, Japan. We sincerely appreciate anonymous reviewers, including those for our previous versions, for their knowledgeable comments. We appreciate Ted Briscoe and Yova Kementchedjhieva for their insightful feedback on the early version of this paper. We also thank the Tohoku NLP Group members, especially Kentaro Inui, for their constructive comments on our earlier work.

References

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel

- Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. [Flamingo: a visual language model for few-shot learning](#). *Proceedings of NeurIPS 2022*, 35:23716–23736.
- Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. 2022. [BEit: BERT pre-training of image transformers](#). In *Proceedings of ICLR 2022*.
- Lawrence W Barsalou. 2008. [Grounded cognition](#). *Annu. Rev. Psychol.*, 59(1):617–645.
- Emily M Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of ACL 2020*, pages 5185–5198.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Proceedings of NeurIPS 2020*.
- Noam Chomsky. 1964. [Aspects of the theory of syntax](#). Technical report, MASSACHUSETTS INST OF TECH CAMBRIDGE RESEARCH LAB OF ELECTRONICS.
- Noam Chomsky. 1980. Rules and representations. *Behavioral and Brain Sciences*, 3(1):1–15.
- 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](#). *arXiv [cs.CL]*.
- Stephen Crain and Mineharu Nakayama. 1987. [Structure dependence in grammar formation](#). *Language*, 63(3):522–543.
- Ekin Dogus Cubuk, Barret Zoph, Jon Shlens, and Quoc Le. 2020. [Randaugment: Practical automated data augmentation with a reduced search space](#). In *Proceedings of NeurIPS 2020*, volume 33, pages 18613–18624.
- 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 *Proceedings of ICLR 2021*.
- Emmanuel Dupoux. 2018. [Cognitive science in the era of artificial intelligence: A roadmap for reverse-engineering the infant language-learner](#). *Cognition*, 173:43–59.
- Jeffrey L Elman. 1990. Finding structure in time. *Cogn. Sci.*, 14(2):179–211.
- 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](#). *Am. J. Speech. Lang. Pathol.*, 26(2):248–265.
- Lila R Gleitman and Henry Gleitman. 1992. [A picture is worth a thousand words, but that’s the problem: The role of syntax in vocabulary acquisition](#). *Current Directions in Psychological Science*, 1(1):31–35.
- Adele Goldberg. 2005. *Constructions at Work: The Nature of Generalization in Language*. Walter de Gruyter GmbH & Co. KG.
- Margaret S Hill and Stacy A Wagovich. 2020. [Word learning from context in school-age children: relations with language ability and executive function](#). *J. Child Lang.*, 47(5):1006–1029.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spacy: Industrial-strength natural language processing in python](#).
- Mark Johnson, Katherine Demuth, and Michael Frank. 2012. [Exploiting social information in grounded language learning via grammatical reduction](#). In *Proceedings of ACL 2012*, pages 883–891.
- Christo Kirov and Ryan Cotterell. 2018. [Recurrent neural networks in linguistic theory: Revisiting pinker and prince \(1988\) and the past tense debate](#). *TACL*, 6:651–665.
- Noriyuki Kojima, Hadar Averbuch-Elor, Alexander Rush, and Yoav Artzi. 2020. [What is learned in visually grounded neural syntax acquisition](#). In *Proceedings of ACL 2020*, pages 2615–2635.
- Julie Anne Legate and Charles D Yang. 2002. [Empirical re-assessment of stimulus poverty arguments](#). *The Linguistic Review*, 19(1-2):151–162.
- Liu, Lin, Cao, Hu, Wei, Zhang, Lin, and Guo. 2021. [Swin transformer: Hierarchical vision transformer using shifted windows](#). In *Proceedings of ICCV 2021*, pages 9992–10002.
- Ilya Loshchilov and Frank Hutter. 2018. [Decoupled weight decay regularization](#). In *Proceedings of ICLR 2018*.
- Chunpeng Ma, Aili Shen, Hiyori Yoshikawa, Tomoya Iwakura, Daniel Beck, and Timothy Baldwin. 2021. [On the \(in\)effectiveness of images for text classification](#). In *Proceedings of EACL 2021*, pages 42–48.
- Ziqiao Ma, Jiayi Pan, and Joyce Chai. 2023. [World-to-words: Grounded open vocabulary acquisition through fast mapping in vision-language models](#). In *Proceedings of ACL 2023*, pages 524–544.

- R Thomas McCoy, Robert Frank, and Tal Linzen. 2018. [Revisiting the poverty of the stimulus: hierarchical generalization without a hierarchical bias in recurrent neural networks](#). In *40th Annual Meeting of the Cognitive Science Society: Changing Minds, CogSci 2018*, pages 2096–2101.
- R Thomas McCoy, Robert Frank, and Tal Linzen. 2020. [Does syntax need to grow on trees? sources of hierarchical inductive bias in Sequence-to-Sequence networks](#). *TACL*, 8:125–140.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of ACL 2019*, pages 3428–3448.
- Colleen McDonough, Lulu Song, Kathy Hirsh-Pasek, Roberta Michnick Golinkoff, and Robert Lannon. 2011. [An image is worth a thousand words: why nouns tend to dominate verbs in early word learning](#). *Dev. Sci.*, 14(2):181–189.
- Tom M Mitchell. 1980. *The need for biases in learning generalizations*. Citeseer.
- Aaron Mueller, Robert Frank, Tal Linzen, Luheng Wang, and Sebastian Schuster. 2022. [Coloring the blank slate: Pre-training imparts a hierarchical inductive bias to sequence-to-sequence models](#). In *Findings of ACL 2022*, pages 1352–1368.
- Mitja Nikolaus, Mostafa Abdou, Matthew Lamm, Rahul Aralikatte, and Desmond Elliott. 2019. [Compositional generalization in image captioning](#). In *Proceedings of CoNLL 2019*, pages 87–98.
- OpenAI. 2023. [Gpt-4 technical report](#). Technical report, OpenAI.
- Amy Perfors, Joshua B Tenenbaum, and Terry Regier. 2011. [The learnability of abstract syntactic principles](#). *Cognition*, 118(3):306–338.
- Laurel Perkins and Jeffrey Lidz. 2021. [Eighteen-month-old infants represent nonlocal syntactic dependencies](#). *Proceedings of the National Academy of Sciences*, 118(41):e2026469118.
- Yulu Qin, Wentao Wang, and Brenden M Lake. 2024. [A systematic investigation of learnability from single child linguistic input](#). *arXiv [cs.CL]*.
- Shaolin Qu and Joyce Chai. 2008. [Incorporating temporal and semantic information with eye gaze for automatic word acquisition in multimodal conversational systems](#). In *Proceedings of EMNLP 2008*, pages 244–253.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of ICML*, pages 8748–8763.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language Models are Unsupervised Multitask Learners](#). Technical report, OpenAI.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. 2021. [Zero-shot text-to-image generation](#). In *Proceedings of ICML 2021*, volume 139, pages 8821–8831.
- Okko Räsänen and Khazar Khorrami. 2019. [A computational model of early language acquisition from audiovisual experiences of young infants](#). In *InterSpeech*, pages 3594–3598. International Speech Communication Association ISCA.
- Deb Roy and Ehud Reiter. 2005. [Connecting language to the world](#). *Artificial Intelligence*, 167(1-2):1–12.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. [Conceptual captions: A cleaned, hypernamed, image alt-text dataset for automatic image captioning](#). In *Proceedings of ACL 2018*, pages 2556–2565.
- Haoyue Shi, Jiayuan Mao, Kevin Gimpel, and Karen Livescu. 2019. [Visually grounded neural syntax acquisition](#). In *Proceedings of ACL 2019*, pages 1842–1861.
- Benjamin Spector. 2007. [Aspects of the pragmatics of plural morphology: On Higher-Order implicatures](#). *Presupposition and Implicature in Compositional Semantics*, pages 243–281.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. 2021. [Training data-efficient image transformers & distillation through attention](#). In *Proceedings of ICML 2021*, volume 139, pages 10347–10357.
- Wai Keen Vong, Wentao Wang, A Emin Orhan, and Brenden M Lake. 2024. [Grounded language acquisition through the eyes and ears of a single child](#). *Science*, 383(6682):504–511.
- Wentao Wang, Wai Keen Vong, Najoung Kim, and Brenden M Lake. 2023. [Finding structure in one child’s linguistic experience](#). *Cogn. Sci.*, 47(6):e13305.
- Alex Warstadt and Samuel R Bowman. 2020. [Can neural networks acquire a structural bias from raw linguistic data?](#) In *Proceedings of Cogsci*, pages 1737–1743.
- Alex Warstadt and Samuel R Bowman. 2022. [What artificial neural networks can tell us about human language acquisition](#). *Algebraic Structures in Natural Language*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel Bowman. 2020a. [BLiMP: The benchmark of linguistic minimal pairs for english](#). *TACL*, 8:377–392.

- 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 EMNLP*, pages 217–235, Online.
- Colin Wilson. 2006. Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cogn. Sci.*, 30(5):945–982.
- Aditya Yedetore, Tal Linzen, Robert Frank, and R Thomas McCoy. 2023. [How poor is the stimulus? evaluating hierarchical generalization in neural networks trained on child-directed speech](#). In *Proceedings of ACL 2023*, pages 9370–9393.
- Tian Yun, Chen Sun, and Ellie Pavlick. 2021. [Does vision-and-language pretraining improve lexical grounding?](#) In *Proceedings of Findings of EMNLP 2021*, pages 4357–4366.
- Chengxu Zhuang, Evelina Fedorenko, and Jacob Andreas. 2024. [Visual grounding helps learn word meanings in low-data regimes](#). In *Proceedings of ACL 2024*, pages 1311–1329.
- Eytan Zweig. 2009. Number-neutral bare plurals and the multiplicity implicature. *Linguistic Philosophy*, 32(4):353–407.

Appendix

A Artificial data

Table 4 shows the textual and visual features used in the ARTIFICIAL dataset. The NUM2 COLOR2 SHAPE2 objects are placed on top of each NUM1 COLOR1 SHAPE1 object, and the VP object is overlaid on the NUM1 COLOR1 SHAPE1 object. We created $3 \times 3 \times 5 \times 4 \times 4 \times 10 = 28,800$ image-caption pairs; 15,000 data are used for training, 1,000 data are used for validation, and 5,000 data are used for evaluation (we sampled 21,000 instances from the 28,800 data).

Category	Word	Visual feature
NUM1/2	a	■
	two	■ ■
	three	■ ■ ■
COLOR1/2	black	■
	red	■
	blue	■
	yellow	■
	lime	■
SHAPE1/2	circle(s)	●
	rectangle(s)	■
	triangle(s)	▲
	hexagon(s)	⬡
VP	walk(s)	👣
	sleep(s)	💤
	run(s) fast	🏃
	wave(s) its hand	👋
	write(s) a text	✍️
	take(s) a bus	🚌
	take(s) a photo	📷
	play(s) soccer	⚽
	play(s) baseball	⚾
throw(s) an arrow at a target	🎯	

Table 4: Vocabularies and their corresponding visual features used in the ARTIFICIAL dataset.

B Vision encoders

All the encoders we used are available in Huggingface. These are pre-trained/fine-tuned on the ImageNet-21k(22k) data with 224^2 resolution and batch size of 16. Table 6 shows the common hyperparameters across the models; other encoder hyperparameters follow the original pre-trained model. To avoid over-fitting, we applied RandAugment (Cubuk et al., 2020) to the

input image and replaced the input image with a white noise with a probability of 0.2. Table 7 shows the image-captioning performance of each model in the validation split of NATURAL data.⁷ The ROUGE score is computed using the implementation of <https://huggingface.co/spaces/evaluate-metric/rouge>. The exact pre-trained models we used are as follows:

Vit:

- <https://huggingface.co/google/vit-base-patch16-224-in21k>
- <https://huggingface.co/google/vit-large-patch16-224-in21k>
- <https://huggingface.co/google/vit-huge-patch14-224-in21k>

Beit:

- <https://huggingface.co/microsoft/beit-base-patch16-224-pt22k-ft22k>
- <https://huggingface.co/microsoft/beit-large-patch16-224-pt22k-ft22k>

Deit:

- <https://huggingface.co/facebook/deit-base-distilled-patch16-224>
- <https://huggingface.co/facebook/deit-small-distilled-patch16-224>
- <https://huggingface.co/facebook/deit-tiny-distilled-patch16-224>

Swin:

- <https://huggingface.co/microsoft/swin-base-patch4-window7-224-in22k>
- <https://huggingface.co/microsoft/swin-large-patch4-window12-384-in22k>

C Evaluation on BLiMP benchmark

We evaluate linguistic knowledge in models with/without vision using the BLiMP benchmark, which has several “circuits” targeting specific linguistic knowledge. Each instance in the circuit is a minimally different sentence pair regarding the targeted grammar item. Similar to our experiment,

⁷Hold-out 1000 AMBIGUOUS instances that do not overlap with the training data.







Vision	OVERALL	ANA. AGR	D-N AGR	IRREGULAR	S-V AGR	ARG. STR	ELLIPSIS	FILLER-GAP	ISLAND	NPI	QUANTIFIERS	BINDING	CTRL. RAIS.
	59.1	60.3	61.1	70.8	61.1	61.0	48.9	65.7	56.5	43.2	72.6	61.4	59.1
	59.1	60.5	60.3	69.3	62.0	61.6	48.9	65.3	55.3	42.5	73.7	62.9	59.5
	58.8	59.4	60.4	70.7	62.3	62.7	42.4	65.7	59.7	42.5	69.8	66.0	61.1

Table 5: Accuracy on each circuit on the BLiMP benchmark. The  model corresponds to the Vit-base model used in the main experiment, the  model corresponds to the model trained with a white noise image, and the  model corresponds to the model trained with shuffled image-caption data.

Decoder	Following the settings in https://huggingface.co/gpt2/blob/main/config.json
Dropout rate in encoder	0.1 (attention and hidden state)
Optimizer	AdamW (Loshchilov and Hutter, 2018)
learning rate	1e-4
betas	(0.9, 0.999)
epsilon	1e-8
Learning scheduler	linear decay
max steps	10,000 (NATURAL setting), 1000 (ARTIFICIAL setting)
warm up steps	0
weight decay	0
Batchsize	512
Beam size	4 (when computing ROUGE)

Table 6: Common hyperparameters across the models with different vision encoders.

we observed whether a model could assign a lower perplexity⁸ to the grammatically correct sentence.









BLiMP has only text input; thus, we must input a sentence alone (and a white noise image) to vision-language models. When inputting only text, a model without vision  might be unfairly favored over a model with vision  from the perspective of the training–inference gap. To achieve a fairer comparison, we also introduce another baseline without proper visual grounding  that is trained with *randomly shuffled* image–caption pairs. We intend that  and  models suffer from a similar degree of handicap regarding the training–inference gap.

Table 5 shows accuracies on each circuit of BLiMP. Vit-base encoder models were evaluated, which are trained using the training set of NATURAL data with 10,000 parameter updates. The

model with vision  does not show a substantial advantage over  and  baselines; this implies that visual input alone cannot enhance their linguistic knowledge.

⁸Sentence pairs in the BLiMP sometimes have different lengths; thus, we avoid using a vanilla probability.

Models	Vis.	NATURAL			ARTIFICIAL		ImageNet
		ROUGE-L F1			ROUGE-L F1		Acc@1
		1,000	5,000	10,000	100	500	
Vit-base (86M)	✓ Δ	32.0 +17.3	35.5 +20.2	37.8 +22.8	80.5 +45.1	100.0 +64.5	84.0
Vit-large (307M)	✓ Δ	30.8 +16.1	35.1 +20.2	37.9 +22.6	76.3 +40.7	100.0 +64.5	85.2
Vit-huge (632M)	✓ Δ	29.2 +14.9	34.1 +18.8	35.8 +20.5	59.1 +23.8	100.0 +63.9	85.1
Beit-base (86M)	✓ Δ	31.7 +15.9	34.5 +19.2	37.4 +22.1	51.5 +16.5	100.0 +64.6	85.2
Beit-large (307M)	✓ Δ	30.4 +15.7	37.0 +21.8	40.2 +24.9	81.2 +46.0	100.0 +64.8	87.4
Deit-base (86M)	✓ Δ	32.2 +18.5	35.6 +20.4	38.2 +22.9	98.5 +63.0	100.0 +64.4	83.4
Deit-small (22M)	✓ Δ	31.0 +16.3	34.6 +19.6	36.6 +21.2	83.0 +47.7	100.0 +64.6	81.2
Deit-tiny (5M)	✓ Δ	30.1 +15.4	33.7 +18.4	35.4 +20.1	93.2 +58.1	100.0 +64.6	74.5
Swin-base (88M)	✓ Δ	34.3 +19.6	37.6 +22.3	40.7 +25.4	99.3 +64.0	100.0 +64.3	85.2
Swin-large (197M)	✓ Δ	34.5 +19.2	38.3 +23.4	41.7 +26.4	97.6 +62.3	100.0 +64.3	87.3
Scratch (86M)	✓ Δ	13.94 +0.16	23.7 +8.78	24.5 +8.93	37.3 +1.88	65.6 30.3	-
Vit-GPT2 (86M)	✓ Δ	32.4 +17.7	35.3 +20.4	37.4 +22.1	93.3 +57.7	100.0 +64.2	84.0

Table 7: ROUGE-L F1 scores of the models at several checkpoints with different training steps. The scores are multiplied by 100. ImageNet accuracy scores are obtained from their original papers.