

CORRESPONDENCES BETWEEN WORD LEARNING IN CHILDREN AND CAPTIONING MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

For human children as well as machine learning systems, a key challenge in learning a word is linking the word to the visual phenomena it describes. By organizing model output into word categories used to analyze child language learning data, we show a correspondence between word learning in children and the performance of image captioning models. Although captioning models are trained only on standard machine learning data, we find that their performance in producing words from a variety of word categories correlates with the age at which children acquire words from each of those categories. To explain why this correspondence exists, we show that the performance of captioning models is correlated with human judgments of the concreteness of words, suggesting that these models are capturing the complex real-world association between words and visual phenomena.

1 INTRODUCTION

Both humans and machines face the problem of establishing the relationship between visual and linguistic information. In humans, this process is known as word learning, and has been extensively studied by developmental scientists. In machines, linking visual features with words is a key part of several tasks, specifically multimodal vision and language tasks. In this paper, we explore correspondences between the solutions humans and machines have found to this problem, starting with the most widely-studied multimodal task, image captioning, and its analogue, children’s word production.

Developmental scientists have long been interested in understanding how infants and young children learn new words (Brown, 1973; Golinkoff et al., 2000; Wojcik et al., 2022; Bloom, 2002; Quine, 1960), often framing the problem as one of establishing reference between words and their corresponding objects, events, or properties (Markman, 1990; Schwab & Lew-Williams, 2016). While the trajectory of word learning varies across children, there is some consistency in the rates at which different kinds of words are learned (Frank et al., 2021). For example, children learning English (as well as many other languages) tend to learn words describing body parts (such as “eye” or “nose”) earlier than they learn connecting words (such as “and” or “because”). Developmental scientists have looked for predictors of this pattern. For example, words that are more frequent in child-directed speech tended to be learned earlier (Swingley & Humphrey, 2018). However, this investigation has been limited to quantities that can be measured from linguistic input (such as word frequency) or by adults making intuitive judgments about the properties of words (such as a word’s “concreteness” or “abstractness”).

Previous work has not explored how the relationship between a word and the visual phenomena it describes relates to the age at which words are learned. Visual aspects of reference pose a challenge for the child learner because scenes vary in complexity (Quine, 1960) and because the referents of words can be highly variable (e.g., “dog” can refer to both chihuahuas and Bernese mountain dogs). Relatedly, some words and word categories refer to concrete objects (e.g., “cup”), but others do not (e.g., “more” or “fine”); concreteness is known to shape age of acquisition (AoA) (Bergelson & Swingley, 2013; Swingley & Humphrey, 2018). Prior experimental work has begun to understand how visual context supports word learning. For example, young children can track word/object co-occurrence statistics over time to disambiguate the meanings of novel labels in complex visual scenes (Chen et al., 2018; Yurovsky et al., 2013), and infants who see more variable views of objects show more rapid vocabulary growth later on (Slone et al., 2019). Although the ease with which a word can be mapped to a concrete visual referent affects children’s noun learning, developmental

scientists have not formalized how the infant mind may process and represent the statistics of visual scenes and labels. Moreover, little is known about how this process may shape word learning.

In this paper, we investigate whether we can capture the visual difficulty of learning words by examining the performance of image captioning systems. Since these systems need to solve a similar problem to children – producing words to communicate about the visual world – they may face the same difficulties. We look at how well image captioning systems perform for different categories of words (such as animals vs. furniture), and examine how the resulting performance measures correspond to children’s word learning. Our results show – across different tasks and architectures – that the difficulty with which machines learn words in different categories is correlated with the difficulty with which children learn words in those categories, and that including this measure in a regression results in a higher correlation with children’s word learning than just using word frequency. We also show that the performance of the captioning models is correlated with human judgments of the “concreteness” of a word, which is known to correlate with AoA. Captioning models thus provide an automated measure of this subjective quantity.

While human children and deep neural networks for image captioning are presumably quite different kinds of systems, discovering correspondences in their performance suggests that some aspects of the difficulty with which different kinds of words are learned is a consequence of the nature of the problem itself. Just as the statistics of the linguistic input to children play a key role in understanding language acquisition (Montag et al., 2015; Laing & Bergelson, 2020; Bergelson & Aslin, 2017), the statistical relationship between that linguistic input and the world that it describes is significant. Our results demonstrate how improvements in computer vision and language systems offer new opportunities for the scientific study of child development.

2 DATASETS AND MODELS

To investigate correspondences in the learning of children and machines, we need two kinds of data:

1. Child word learning data, including which words (or word categories) are learned at various developmental stages.
2. Standardized image and natural language machine learning data, which can be used to train vision and language models to produce language comparable to child word production.

We address both these needs by working with multiple sources of data: WordBank for child word learning (Frank et al., 2017), and COCO for training our models (Lin et al., 2014).

2.1 CHILD LANGUAGE ACQUISITION DATA: WORDBANK

We use data from the WordBank child language database (Frank et al., 2017) to extract words commonly produced by toddlers between the ages of 16 and 30 months. Figure 1 gives an example of the type of data collected and tracked by WordBank. In particular, the data we use indicates which words are produced by toddlers of various ages. WordBank contains these production percentiles for approximately 1200 words, which we will use in our analysis.

In order to compare word learning at scale, we decided to investigate patterns in word categories instead of individual words. For an effective comparison to child word learning, we use word categories for which there exists parallel child data. Fortunately, the WordBank database contains such data. We extracted approximately 1200 frequently-produced words for toddlers, as listed in the WordBank database, and mapped them to the corresponding category. WordBank categories include people, toys, animals, etc. We remove sounds/sound effects (such as “cockadoodledo”) from these categories, because our models are restricted to vision and language modalities.

2.2 COMPUTER VISION TASKS: COCO

On the computer vision side we use the canonical COCO (Lin et al., 2014) image captioning dataset (the Karpathy & Fei-Fei (2015) split, to be precise) for training and evaluating the models. The dataset contains 113,278 training and 5,000 validation images, each associated with 5 captions provided by human annotators. Our main experiments involve image captioning, where a model is trained




Figure 1: WordBank data tracking toddlers’ production of different categories across different ages. Each line represents the average proportion of children producing all words in a given category on the MCDI, and the age at which half of children produce a given word (age of acquisition; AoA) is marked with a circle on the trajectory of each word category. There is a consistent ordering in the acquisition of words from different categories, which we compare to computer vision models.

to produce a natural language caption. To provide a baseline and to validate our findings, we also present results on a simpler image classification task, where a model is trained to predict which visual categories are present in the image without tying these categories to a natural language description. Although image classification cannot fully capture the complexities of words contextualized in natural language, it is a simpler task that requires linking visual cues to individual words, and as such a useful baseline to compare to captioning. For classification, we create a multi-label, multi-class, setup with binary labels comprising of 855 individual words which are in both COCO and WordBank. The binary label indicates whether the word is mentioned in any of the captions associated with the image.

We run experiments with canonical captioning and classification models. Our goal is to verify that our findings hold across a range of standard setups. For image captioning, we explore two standard vision backbones: a ResNet101 CNN (He et al., 2016) or bottom-up features from Faster R-CNN (Ren et al., 2015; Anderson et al., 2018). We combine these models with one of two language models: the classic LSTM (Anderson et al., 2018) or the more recent Transformer (Vaswani et al., 2017). For the image classification baseline models we use two standard CNN architectures: VGGNet (Simonyan & Zisserman, 2014) and ResNet50 (He et al., 2016), both with and without pretraining on ImageNet (Russakovsky et al., 2015).

We use open-source implementations of all models along with the recommended hyperparameters (Luo et al., 2018; 2021; Chollet et al., 2015). Each model is trained on a single GPU. For the LSTM layers of the captioning models, we use an input encoding size of 1000, a hidden size of 512, a batch size of 10, and an adaptive learning rate. For the Transformer layers, we use 8 attention heads, 6 encoder and decoder layers, 512 hidden unit size, and a batch size of 10. For image classification, we train with an adaptive learning rate using the Adam (Kingma & Ba, 2014) optimizer and dropout until the loss converges. VGGNet trained from scratch proved difficult to train to convergence, despite performing grid search over initial learning rate, dropout, and batch size, so only pretrained results are reported. We keep as much consistent as possible between implementations to allow comparison.

3 METRICS

To quantitatively measure the correspondence between word learning in human children and computer vision systems, we used standard metrics used in each field: the median age at which children produce a word, and the word-level performance measures of SPICE for captioning and AUC for classification.

3.1 METRIC FOR CHILDREN: AOA (AGE OF ACQUISITION)

The Age of Acquisition (AoA) of a word is defined as the age at which at least 50% of children produce the word. WordBank includes the vocabularies of 5520 toddlers learning North American

English assessed using parent report on the MacArthur Bates Communicative Development Inventory (MCDI) (Fenson et al., 2007). In the WordBank database, AoA can be calculated over the parent-reported scores for word learning for each child in the database. AoA has previously been shown to correspond with the difficulty of learning to read a word (Coltheart et al., 1988), and we use it here to quantify the difficulty of learning a word for a child. AoA was calculated for each word and then averaged within each of the WordBank categories: body parts, animals, vehicles, toys, household, outside, food/drink, furniture/rooms, clothing, locations, descriptive words, places, people, action words, pronouns, question words, quantifiers, helping verbs, time words, and connecting words.

3.2 METRICS FOR MACHINES: SPICE AND AUC

The metrics we use for our models measure performance at the level of individual words. For captioning models, we use the Semantic-Propositional Image Caption Evaluation score (SPICE) (Anderson et al., 2016). SPICE is an automatic evaluation metric which uses scene graphs corresponding to the actual image to gauge semantic and propositional correctness, instead of just the textual n-gram comparison of previous metrics.

From a caption like “woman sitting on a brown chair in a restaurant,” SPICE produces a tuple-based scene graph containing tuples like “woman-sitting” and “sitting-on-chair.” SPICE then calculates whether each produced tuple matches one of the tuples that appear in the ground truth manual captions. To get a score for each individual WordBank word, we then average the tuple-based scores of all the tuples where the word appears. The intuition for using this metric is that it is impossible to gauge whether a word like “sitting” is used correctly without looking at the other words around it.

For the classification baseline models, we use AUC (the area under the receiver operating characteristic (ROC) curve; Freedman (2009)) as a classification metric which is robust to class imbalance. Our multi-label, multi-class classification task was binary for each label, so a per-word AUC calculated over each label was an appropriate metric.

For both SPICE and AUC, after calculating at a word-level (or a tuple-level, for SPICE) we then aggregate over WordBank categories for ease of comparison to AoA for those same categories.

4 RESULTS

As an initial analysis, we examined the raw correlation between AoA and the machine metrics. We then conducted multiple regression analyses to see whether computer vision systems can improve correlations with AoA over existing measures used in the child language acquisition literature.

4.1 CORRELATIONS

To compare category-level AoA to SPICE/AUC, we report two types of correlation: the Pearson correlation coefficient, which assumes a linear relationship, and the Spearman rank-order correlation coefficient, which only assumes monotonicity Freedman (2009). The results are shown in Table 2 for classification and Table 1 for captioning, and the corresponding scatterplots appear in Figure 2.

Table 1: Captioning model SPICE correlation with AoA

Captioning architecture		Pearson	p	Spearman	p
Pretrained Visual Features	Language Layers				
ResNet 101	LSTM	-0.515	0.034	-0.640	0.006
Bottom-up (Faster R-CNN)	LSTM	-0.617	0.004	-0.708	0.000
Bottom-up (Faster R-CNN)	Transformer	-0.565	0.012	-0.624	0.004

The three captioning models showed statistically significant correlations with AoA. One classification baseline model (ResNet50 with pretrained features) also produced a statistically significant correlation. In all four of these models, as performance (SPICE or AUC) increased, AoA decreased: categories of words that were easier for the models were acquired earlier by children. The remaining two classification models (VGG with pretrained features and ResNet50 trained from scratch) also showed

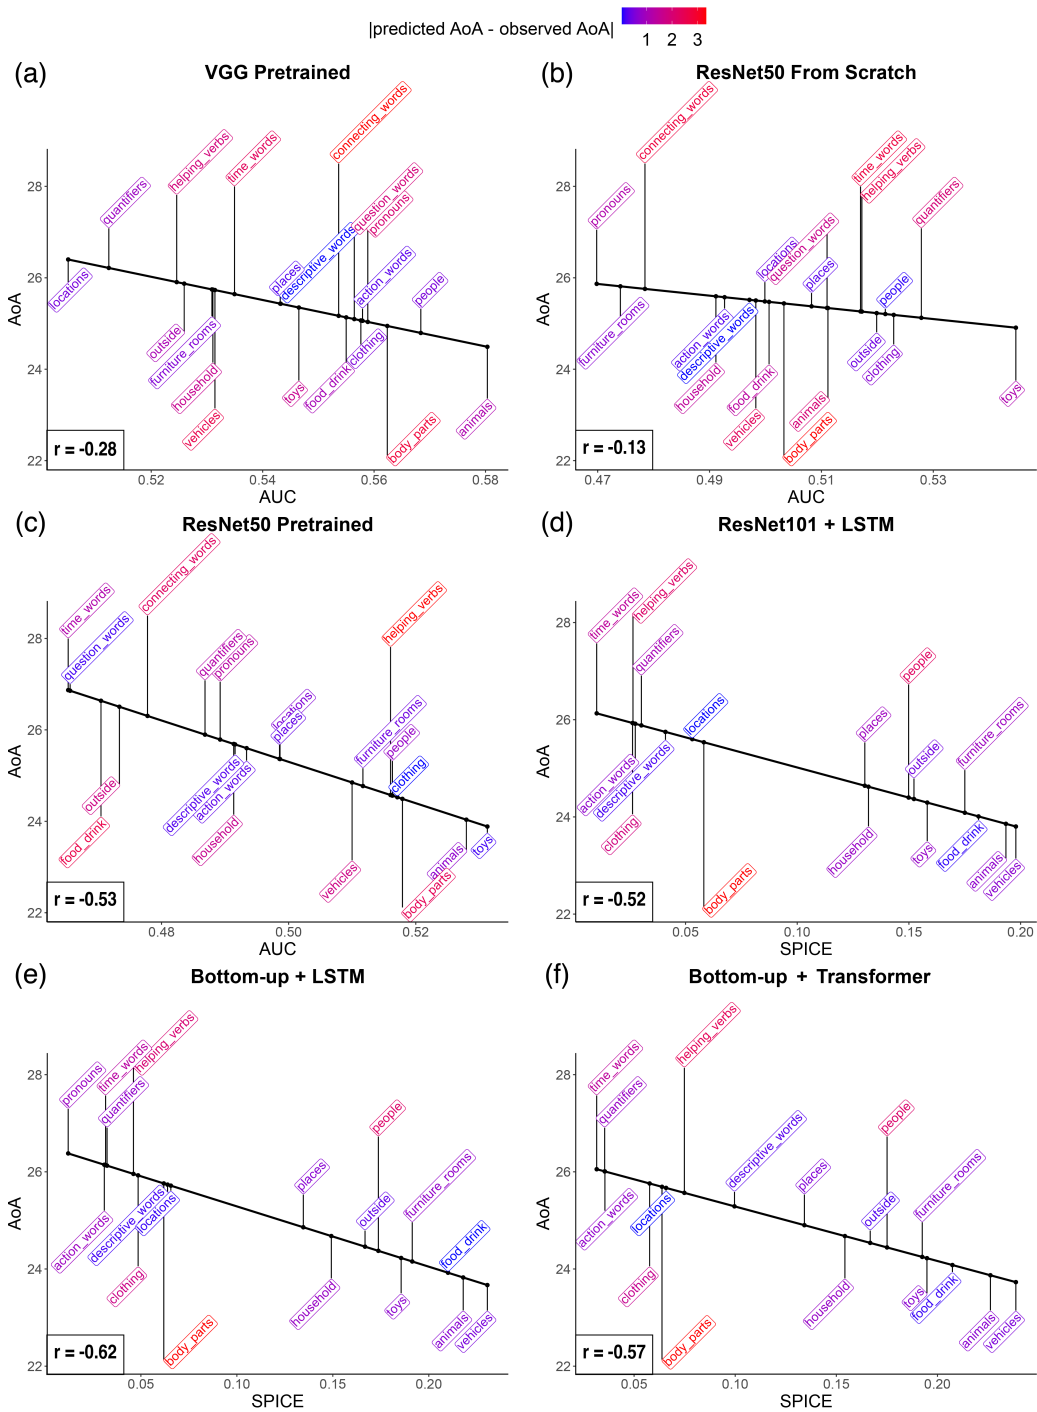


Figure 2: Regression of successful models’ performance vs age of acquisition in months, per category. Each category label is placed so that the corner from which the vertical line emerges is the corresponding age of acquisition (AoA) and AUC (classification) or SPICE (captioning) values. The black regression line shows the AoA produced by each model in the regression, and the vertical residual lines and category labels show the observed average AoA for the category.

Table 2: Classification model AUC correlation with AoA

Classification model setup					
Model	Training	Pearson	p	Spearman	p
VGG	Pretrained	-0.280	0.232	-0.311	0.182
ResNet50	From scratch	-0.138	0.562	-0.081	0.734
ResNet50	Pretrained	-0.531	0.016	-0.544	0.013

correlations consistent with this relationship, but those correlations were not statistically significant. The correlation for ResNet50 trained from scratch was particularly weak, suggesting that pretrained features may be important. The correlations for the captioning systems were all of similar magnitude, suggesting that the specific architecture (including the choice of an LSTM or Transformer) may be less relevant than the captioning (natural language) task itself.

4.2 COMPARISON WITH OTHER MEASURES

As noted above, developmental psychologists have explored variables that correlate with children’s word learning. One such variable is the frequency with which words appear in child-directed speech. We evaluated the correlation between word frequency (extracted from the TalkBank database; MacWhinney (2007)) and AoA, finding a Pearson correlation of $r = -0.377$ ($p = 0.092$) and a statistically significant Spearman correlation of $\rho = -0.494$ ($p = 0.022$). Notably, this correlation is of similar magnitude to those for the metrics from our models. See Appendix for a scatterplot.

To determine whether our models’ metrics (SPICE and AUC) correlate with AoA even after accounting for word frequency, we conducted a multiple regression analysis where the independent variables are word frequency and SPICE/AUC, and the dependent variable is AoA. The coefficients of the multiple regression analysis show that SPICE/AUC across different successful models do indeed correlate with AoA over and above frequency in child-directed data. The results are shown in Table 3. The multiple regression analysis showed that all four measures that originally resulted in a statistically significant correlation with AoA remained statistically significant when word frequency was taken into account. Notably, word frequency was no longer statistically significant in the resulting models.

Table 3: Multiple regression with TalkBank word frequency in addition to AUC or SPICE

Model	Regression Coefficients (with AoA)				
	TalkBank frequency	p	SPICE / AUC	p	R^2
ResNet101 + LSTM	-1.297	0.267	-2.148	0.050	0.329
Bottom-up + LSTM	-1.225	0.229	-2.769	0.007	0.433
Bottom-up + Transformer	-1.552	0.153	-2.355	0.027	0.404
VGG Pretrained	-1.891	0.128	-1.580	0.306	0.199
ResNet50 From Scratch	-2.160	0.087	-1.243	0.431	0.178
ResNet50 Pretrained	-1.286	0.269	-2.559	0.044	0.333

Another variable that has been shown to be correlated with AoA is the “concreteness” of words Swingley & Humphrey (2018); Bergelson & Swingley (2013). Unlike word frequency, concreteness is not a property that can be measured directly from the linguistic input to children. Rather, it is typically measured by asking human raters to rate on a scale how “concrete” or “abstract” they consider a word to be, after providing a definition for concrete (e.g., can be experienced with the five senses) and abstract words (e.g., cannot be experienced through the five senses) Brysbaert et al. (2014). Some work has expanded these rating lists by using supervised models trained directly to predict concreteness Köper & im Walde (2017). We ran a second multiple regression including a standard measure of concreteness (taken from Köper & im Walde (2017)). The results are shown in Table 6. In this model, concreteness was the only statistically significant variable, with neither word frequency nor our measures being statistically significant. The same result is observed in a multiple regression incorporating only concreteness and SPICE/AUC 4.

Table 4: Multiple regression with concreteness judgments in addition to AUC or SPICE

Model	Regression Coefficients (with AoA)				R^2
	SPICE / AUC	p	Concreteness	p	
ResNet101 + LSTM	0.620	0.574	-3.996	0.002	0.631
Bottom-up + LSTM	0.488	0.668	-4.024	0.002	0.652
Bottom-up + Transformer	0.544	0.621	-3.897	0.002	0.629
VGG Pretrained	-0.128	0.884	-3.925	0.000	0.754
ResNet50 From Scratch	-0.409	0.631	-3.926	0.000	0.757
ResNet50 Pretrained	-0.940	0.213	-3.601	0.000	0.776

Table 5: Correlation of SPICE/AUC with human judgments of concreteness

Model	Pearson	p	Spearman	p
ResNet101 + LSTM	0.733	0.001	0.598	0.011
Bottom-up + LSTM	0.744	0.000	0.659	0.003
Bottom-up + Transformer	0.690	0.002	0.574	0.016
VGG Pretrained	0.303	0.195	0.277	0.238
ResNet50 From Scratch	0.092	0.701	0.056	0.816
ResNet50 Pretrained	0.459	0.042	0.421	0.064

Investigation of this result revealed that it is a consequence of substantial collinearity between the computer vision measures and concreteness ratings. The correlations between SPICE/AUC and concreteness are shown in Table 5. SPICE scores from all three captioning models correlate very strongly with human judgements of concreteness. AUC from the pretrained ResNet50 classification baseline has a weaker, but still statistically significant correlation to concreteness.

The observed correlation between concreteness and SPICE makes sense: concreteness is people’s judgment of how well a word corresponds with a visible or tangible thing in the world, and this is what our measures reflect. A high correlation with concreteness thus indicates that our captioning models are capturing what we intended: the ease of relating a word to its visual referent(s). Importantly, the performance of a captioning model is an objective quantity that can be estimated directly from a dataset, rather than a subjective quantity that requires additional judgments from people. Further, our approach captures meaningful variability within the visual contexts of concrete and abstract words. For example, the words “hello” and “economy” may be judged as equally abstract (Köper & im Walde, 2017), however, “hello” may be associated with a more consistent visual context as part of a routine (e.g., waving) compared to “economy”. Similarly, concrete words like “spoon” may have a more consistent surrounding visual context (e.g., a kitchen), compared to words like “dog”, which may be encountered in many different visual contexts. Future work can apply this approach to go beyond category-level estimates and capture the visual variability of different items, as well as individual differences in the visual contexts that different children experience in densely sampled child-view visual corpora (e.g., Sullivan et al. (2022)). By providing a new way to directly measure the concreteness of words, our approach provides a novel metric that can be used in the psychological investigation of language processing more broadly.

5 RELATIONSHIP TO PREVIOUS WORK

While there has been no previous work looking directly at the relationship of AoA with computer vision models, there is an extensive literature in cognitive science and computer vision examining correspondences between humans and machines. For example, representations from image classification systems have been used to predict human judgments of image typicality (Lake et al., 2015), the similarity between images (Jozwik et al., 2017; Peterson et al., 2018; Hebart et al., 2020), image classification (Sanders & Nosofsky, 2020; Battleday et al., 2020), and neural responses to images (Yamins & DiCarlo, 2016; Schrimpf et al., 2020). Better capturing these aspects of human

Table 6: Multiple regression with all variables: TalkBank word frequency, SPICE/AUC, and concreteness

Model	Regression Coefficients (with AoA)						
	TalkBank frequency	p	SPICE / AUC	p	Concreteness	p	R^2
ResNet101 + LSTM	-1.469	0.076	0.901	0.383	-4.099	0.001	0.713
Bottom-up + LSTM	-1.239	0.126	0.647	0.552	-4.106	0.001	0.707
Bottom-up + Transformer	-1.445	0.082	0.782	0.447	-3.951	0.001	0.709
VGG Pretrained	-1.111	0.096	-0.015	0.986	-3.747	0.000	0.794
ResNet50 From Scratch	-1.170	0.077	-0.594	0.461	-3.704	0.000	0.801
ResNet50 Pretrained	-0.967	0.144	-0.661	0.374	-3.530	0.000	0.804

cognition has been shown to result in improvements in computer vision systems (Peterson et al., 2019). Developmental research has also explored the use of deep neural networks to capture aspects of children’s language learning, particularly systems that are trained on data from cameras mounted on the heads of infants (Bambach et al., 2018; Orhan et al., 2020; Tsutsui et al., 2020). This research has primarily focused on visual object learning rather than the timecourse of word-learning. Other work has used multimodal neural networks to capture human performance in stylized word-learning settings (Vong & Lake, 2022). This work provides a converging perspective on how models from computer vision can be used to capture the relationship between linguistic and visual input.

6 DISCUSSION

We have shown that despite training on only standard machine learning datasets (COCO and ImageNet), several captioning models and one classification model show correspondences with the age at which children acquire different categories of words. This result holds across multiple captioning architectures and model types. Furthermore, captioning model performance correlates strongly with human judgments of concreteness, more so than even the best classification model. Taken together, these results indicate that captioning models capture part of the visual difficulty of learning a word.

Figure 2 provides some intuition for why visual difficulty goes beyond training data distribution: for example, while the categories “food/drink” and “descriptive words” occur much more frequently in child-directed speech than in the COCO training data, the models correspond well with AoA for those categories. This illustrates the value of ML approaches to concreteness, and provides some intuition for the commonalities in child and model learning. Certain categories are also difficult for both models and children, despite those categories being *over*represented in the training data. For example, quantifiers are difficult for both models and children to learn, despite being well represented in COCO training data and child-directed speech (distributions in the Appendix).

Pretraining seems to be one of the key differentiating factors between the models which showed substantial correlation and those which did not, such as the ResNet50 model pretrained and the same ResNet50 trained from scratch. There are several potential reasons for this. If pretraining (even on ImageNet alone) allows for learning how to extract visual features, that skill can then be applied to more complex visual features and patterns than those in the training data. However, pretraining is not the whole story: the difference between pretrained VGGNet and ResNet50 classification models’ correlation to AoA shows that architecture does contribute to the correlation as well.

The consistently high correlations for captioning models with language components support anecdotal evidence that larger models combining vision and language modalities produce more human-like performance for visual word learning. The high correlation across different architectures opens the door to future investigations as to why this is the case – it is clearly not one particular element, such as a transformer language model, that yields this result. And yet, sophisticated language modeling with attention mechanisms, present in all the captioning models either through the LSTM or Transformer layers, seems to contribute to the robustness of the high correlation for captioning models.

6.1 LIMITATIONS

One limitation of our work, as applied to modeling child behavior, is that we are focused on captioning models trained on a standard data set (COCO) that differs substantially from the linguistic input received by human children. However, in some ways this strengthens our results and renders our findings all the more surprising: a meaningful correspondence emerges *despite* the difference in training data. We would anticipate that training on datasets that are more similar to the linguistic experience of human children (e.g., Sullivan et al. (2022)) would only increase this correspondence.

A second limitation is that we are focused on word categories rather than individual words. The previous literature in developmental science has evaluated correlations with word frequency and concreteness at the level of individual words. We were unable to do so here partly because the mismatch between the distributions of COCO and child-directed speech is more substantial at the level of individual words than at the level of word categories (even at the level of word categories there is little correspondence, as shown in the Appendix). This was one reason we decided to focus on word categories. However, using datasets that have a closer correspondence to child-directed speech will make it possible to perform similar analyses for individual words.

Finally, we have only explored a subset of the vast array of computer vision models and metrics. The subset we selected was intended to provide a representative sampling of different architectures, tasks, and training regimes. However, a more comprehensive survey of the field may be necessary in order to draw strong conclusions about the impact of these choices on the correspondence between word learning in humans and machines. Unlike other recent evaluations of the correspondence between computer vision models and aspects of human cognition, this is not a simple matter of downloading a pretrained model and evaluating its representations: classification models need to be trained on the multi-way classification task we used to make them comparable to captioning models, and adoption of other datasets will require even more extensive retraining of models. This means that a broader evaluation of computer vision models is a significant undertaking.

6.2 FUTURE WORK

In demonstrating how vision and language models effectively capture word learning difficulty in children, this work also opens the door to more behavioral comparisons of word learning in children and multimodal models. We have demonstrated this result on a standardized group of datasets, with standard pretraining protocols. An important future question is, to what extent particular architectural components (ResNet/Faster R-CNN visual features, or LSTM/Transformer language layers) are important for capturing different facets of child word learning. Is it the scale of larger captioning models which yields the robust similarity to child word learning? Or is it the attention mechanisms in the more sophisticated language components? Another important line of inquiry is how this result changes with datasets; it is surprising that this correlation exists although children and models are certainly exposed to different data. Training models on a child-directed dataset, such as SAYCam Sullivan et al. (2022) is likely to strengthen the correlation to child word learning patterns. Our results lay the groundwork for further behavioral comparisons between models and child learning.

6.3 CONCLUSION

We have shown that the difficulty with which image captioning models learn different categories of words corresponds to the age at which children learn words in those categories. Although captioning models and human children potentially have significant differences in the mechanisms of learning, both face the challenge of relating a word to the visual phenomena it describes. The developmental parallels, which show that the difficulty of learning different categories of words is aligned for both machines and children, suggest that the structure of the learning problem itself may induce similarities in patterns of learning. These results already provide useful insights for both cognitive scientists and machine learning researchers: For cognitive scientists, our models provide an approximation of human concreteness judgements, useful for modeling human behavior. For ML researchers, this work provides a blueprint for understanding the behaviors and developmental patterns of multimodal models, and meaningfully comparing them to humans. We hope that these results open the door to new opportunities to model child development using machine learning systems for computer vision and language, and in turn help us to understand these machine learning systems better through their parallels to child development.

REFERENCES

- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *European Conference on Computer Vision*, pp. 382–398. Springer, 2016.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6077–6086, 2018.
- Sven Bambach, David Crandall, Linda Smith, and Chen Yu. Toddler-inspired visual object learning. *Advances in Neural Information Processing Systems*, 31, 2018.
- Ruairidh M Battleday, Joshua C Peterson, and Thomas L Griffiths. Capturing human categorization of natural images by combining deep networks and cognitive models. *Nature Communications*, 11(1):1–14, 2020.
- Elika Bergelson and Richard N Aslin. Nature and origins of the lexicon in 6-mo-olds. *Proceedings of the National Academy of Sciences*, 114(49):12916–12921, 2017.
- Elika Bergelson and Daniel Swingley. The acquisition of abstract words by young infants. *Cognition*, 127(3):391–397, 2013.
- Paul Bloom. *How children learn the meanings of words*. MIT Press, 2002.
- Roger Brown. *A first language: The early stages*. Harvard University Press, 1973.
- Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. Concreteness ratings for 40 thousand generally known english word lemmas. *Behavior Research Methods*, 46(3):904–911, 2014.
- Chi-hsin Chen, Yayun Zhang, and Chen Yu. Learning object names at different hierarchical levels using cross-situational statistics. *Cognitive science*, 42:591–605, 2018.
- François Chollet et al. Keras. <https://keras.io>, 2015.
- Veronika Coltheart, Veronica J Laxon, and Corriene Keating. Effects of word imageability and age of acquisition on children’s reading. *British Journal of Psychology*, 79(1):1–12, 1988.
- Larry Fenson et al. *MacArthur-Bates communicative development inventories*. Paul H. Brookes Publishing Company Baltimore, MD, 2007.
- Michael C Frank, Mika Braginsky, Daniel Yurovsky, and Virginia A Marchman. Wordbank: An open repository for developmental vocabulary data. *Journal of Child Language*, 44(3):677–694, 2017.
- Michael C Frank, Mika Braginsky, Daniel Yurovsky, and Virginia A Marchman. *Variability and consistency in early language learning: The Wordbank project*. MIT Press, 2021.
- David A Freedman. *Statistical models: theory and practice*. Cambridge University Press, 2009.
- Roberta Michnick Golinkoff, Kathryn Hirsh-Pasek, Lois Bloom, Linda B Smith, Amanda L Woodward, Nameera Akhtar, Michael Tomasello, and George Hollich. *Becoming a word learner: A debate on lexical acquisition*. Oxford University Press, 2000.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Martin N Hebart, Charles Y Zheng, Francisco Pereira, and Chris I Baker. Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature Human Behaviour*, 4(11):1173–1185, 2020.
- Kamila M Jozwik, Nikolaus Kriegeskorte, Katherine R Storrs, and Marieke Mur. Deep convolutional neural networks outperform feature-based but not categorical models in explaining object similarity judgments. *Frontiers in Psychology*, 8:1726, 2017.

- Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3128–3137, 2015.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Maximilian Köper and Sabine Schulte im Walde. Improving verb metaphor detection by propagating abstractness to words, phrases and individual senses. In *Proceedings of the 1st Workshop on Sense, Concept and Entity Representations and their Applications*, pp. 24–30, 2017.
- Catherine Laing and Erika Bergelson. From babble to words: Infants’ early productions match words and objects in their environment. *Cognitive Psychology*, 122:101308, 2020.
- Brenden M Lake, Wojciech Zaremba, Rob Fergus, and Todd M Gureckis. Deep neural networks predict category typicality ratings for images. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, 2015.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pp. 740–755. Springer, 2014.
- Ruotian Luo, Brian Price, Scott Cohen, and Gregory Shakhnarovich. Discriminability objective for training descriptive captions. *arXiv preprint arXiv:1803.04376*, 2018.
- Ruotian Luo et al. An image captioning codebase. <https://github.com/ruotianluo/ImageCaptioning.pytorch>, 2021. MIT License.
- Brian MacWhinney. The talkbank project. In *Creating and digitizing language corpora*, pp. 163–180. Springer, 2007.
- Ellen M Markman. Constraints children place on word meanings. *Cognitive Science*, 14(1):57–77, 1990.
- Jessica L Montag, Michael N Jones, and Linda B Smith. The words children hear: Picture books and the statistics for language learning. *Psychological Science*, 26(9):1489–1496, 2015.
- Emin Orhan, Vaibhav Gupta, and Brenden M Lake. Self-supervised learning through the eyes of a child. *Advances in Neural Information Processing Systems*, 33:9960–9971, 2020.
- Joshua C Peterson, Joshua T Abbott, and Thomas L Griffiths. Evaluating (and improving) the correspondence between deep neural networks and human representations. *Cognitive Science*, 42(8):2648–2669, 2018.
- Joshua C Peterson, Ruairidh M Battleday, Thomas L Griffiths, and Olga Russakovsky. Human uncertainty makes classification more robust. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9617–9626, 2019.
- Willard Van Orman Quine. *Word and Object*. MIT Press, 1960.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28, 2015.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- Craig A Sanders and Robert M Nosofsky. Training deep networks to construct a psychological feature space for a natural-object category domain. *Computational Brain & Behavior*, 3(3):229–251, 2020.
- Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J Majaj, Rishi Rajalingham, Elias B Issa, Kohitij Kar, Pouya Bashivan, Jonathan Prescott-Roy, Franziska Geiger, et al. Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, pp. 407007, 2020.

- Jessica F Schwab and Casey Lew-Williams. Repetition across successive sentences facilitates young children’s word learning. *Developmental Psychology*, 52(6):879, 2016.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Lauren K Slone, Linda B Smith, and Chen Yu. Self-generated variability in object images predicts vocabulary growth. *Developmental Science*, 22(6):e12816, 2019.
- Jessica Sullivan, Michelle Mei, Andrew Perfors, Erica Wojcik, and Michael C Frank. SAYCam: A large, longitudinal audiovisual dataset recorded from the infant’s perspective. *Open Mind*, 5:20–29, 2022.
- Daniel Swingley and Colman Humphrey. Quantitative linguistic predictors of infants’ learning of specific english words. *Child Development*, 89(4):1247–1267, 2018.
- Satoshi Tsutsui, Arjun Chandrasekaran, Md Alimoor Reza, David Crandall, and Chen Yu. A computational model of early word learning from the infant’s point of view. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017.
- Wai Keen Vong and Brenden M Lake. Cross-situational word learning with multimodal neural networks. *Cognitive Science*, 46(4):e13122, 2022.
- Erica H Wojcik, Martin Zettersten, and Viridiana L Benitez. The map trap: Why and how word learning research should move beyond mapping. *Wiley Interdisciplinary Reviews: Cognitive Science*, pp. e1596, 2022.
- Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3):356–365, 2016.
- Daniel Yurovsky, Linda B Smith, and Chen Yu. Statistical word learning at scale: The baby’s view is better. *Developmental Science*, 16(6):959–966, 2013.

A APPENDIX




Figure 3: Distribution of word categories in the models’ training data (COCO) compared to the TalkBank child-directed data.

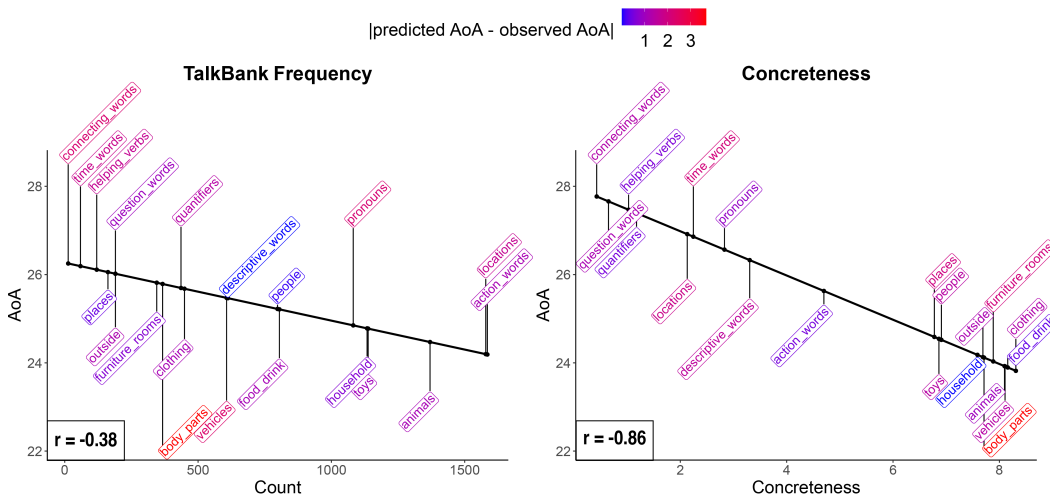


Figure 4: Regression results for word frequency in child-directed TalkBank data, as well as concreteness judgements, each individually correlated with AoA per category.