

Caregiver Talk Shapes Toddler Vision: A Computational Study of Dyadic Play

Timothy Schaumlöffel^{1,2,*}, Arthur Aubret^{3,4,*}, Gemma Roig^{1,2,†}, Jochen Triesch^{1,3,†}

Abstract—Infants’ ability to recognize and categorize objects develops gradually. The second year of life is marked by both the emergence of more semantic visual representations and a better understanding of word meaning. This suggests that language input may play an important role in shaping visual representations. However, even in suitable contexts for word learning like dyadic play sessions, caregivers utterances are sparse and ambiguous, often referring to objects that are different from the one to which the child attends. Here, we systematically investigate to what extent caregivers’ utterances can nevertheless enhance visual representations. For this we propose a computational model of visual representation learning during dyadic play. We introduce a synthetic dataset of ego-centric images perceived by a toddler-agent that moves and rotates toy objects in different parts of its home environment while “hearing” caregivers’ utterances, modeled as captions. We propose to model toddlers’ learning as simultaneously aligning representations for 1) close-in-time images and 2) co-occurring images and utterances. We show that utterances with statistics matching those of real caregivers give rise to representations supporting improved category recognition. Our analysis reveals that a small decrease/increase in object-relevant naming frequencies can drastically impact the learned representations. This affects the attention on object names within an utterance, which is required for efficient visuo-linguistic alignment. Overall, our results support the hypothesis that caregivers’ naming utterances can improve toddlers’ visual representations.

I. INTRODUCTION

Second-year toddlers are proficient learners of object representations that support the recognition of objects independently of viewpoint (object instance recognition) and the assignment of novel exemplars to learned categories (object categorization recognition). What learning mechanisms support this development? On one side, there is evidence that biological organisms learn similar representations for close-in-time visual inputs [1], [2]. This is called the *slowness principle* [3] and it supports the building of view-invariant object representations during, e.g., object manipulations [4], [5]. On the other side, the acquisition of object names correlates with a stronger focus on global shape features relative to texture [6], [7], [8], which spurs performance in name-agnostic categorization tasks [9], [10], see also [11]. For learning object names, the synchronization of toddlers’ visual attention on the object and its naming by the caregivers is crucial [12]. Dyadic play could be particularly

important for promoting the emergence of more semantic object representations, because both learning mechanisms are likely to co-occur during such play: the toddler extensively focuses on/manipulates objects and the caregiver frequently talks to the toddler about the present objects [13], [14]. However, naming utterances during such play are *sparse* [15] and *ambiguous*: they typically contain several words unrelated to the object. Furthermore, even though toddlers have biased attention towards held objects [16], they often play in contexts that admit several other objects in the background [17], [18], adding another potential source of confusion.

In this paper, we investigate to what extent dyadic play sessions can support the emergence of semantic object representations in toddlers. We tackle this question by proposing a computational model of toddlers’ learning during dyadic play. First, we introduce the Dyadic Play Dataset, a novel dataset of images simulating at-home ego-centric dyadic play sessions. In this dataset, we simulate a playing toddler that moves and turns 3D toy models extracted from the Toy4k dataset [19]. The agent “plays” in front of a background composed of household furniture (cf. Fig. 1A) and a plausible number of toys scattered on the floor (cf. Fig. 1B). As shown in Fig. 1C, The agent occasionally “hears” caregiver’s utterances, modeled as captions extracted from the CHILDES database, a comprehensive data repository of children’s language acquisition including child-directed speech [20]. Then, we consider a bio-inspired model of toddlers’ learning that 1) maps close-in-time visual inputs to similar representations and 2) similarly aligns the representations of visual inputs and co-occurring naming utterances (Fig. 1D). We further simulate toddlers’ visual attention by biasing the model to extract visual features from the currently held object. The overall model allows us to systematically study the potential impact of the sparsity and ambiguity of naming utterances on the learned visual representations.

Our experiments show that utterance statistics reported in developmental psychology experiments support the construction of semantic visual representations. Our analysis shows that realistic object-relevant naming frequencies fall within a range of values for which a multiplicative factor of two drastically impact the learned representation. This impacts whether the model attends to the object name within an utterance, which is necessary to efficiently guide visual representations. Thus, our paper provides computational support to the hypothesis that caregivers’ sparse and ambiguous utterances help to build visual representations during dyadic play sessions.

*Equal contribution †Shared last authorship

¹Goethe University, Frankfurt am Main, Department of Computer Science, Germany, {lastname}@em.uni-frankfurt.de

²The Hessian Center for Artificial Intelligence (hessian.AI), Darmstadt, Germany

³Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany, {lastname}@fias.uni-frankfurt.de

⁴Xidian-FIAS international Joint Research Center, Ruth-Moufang-Strasse, D-60438 Frankfurt am Main, Germany

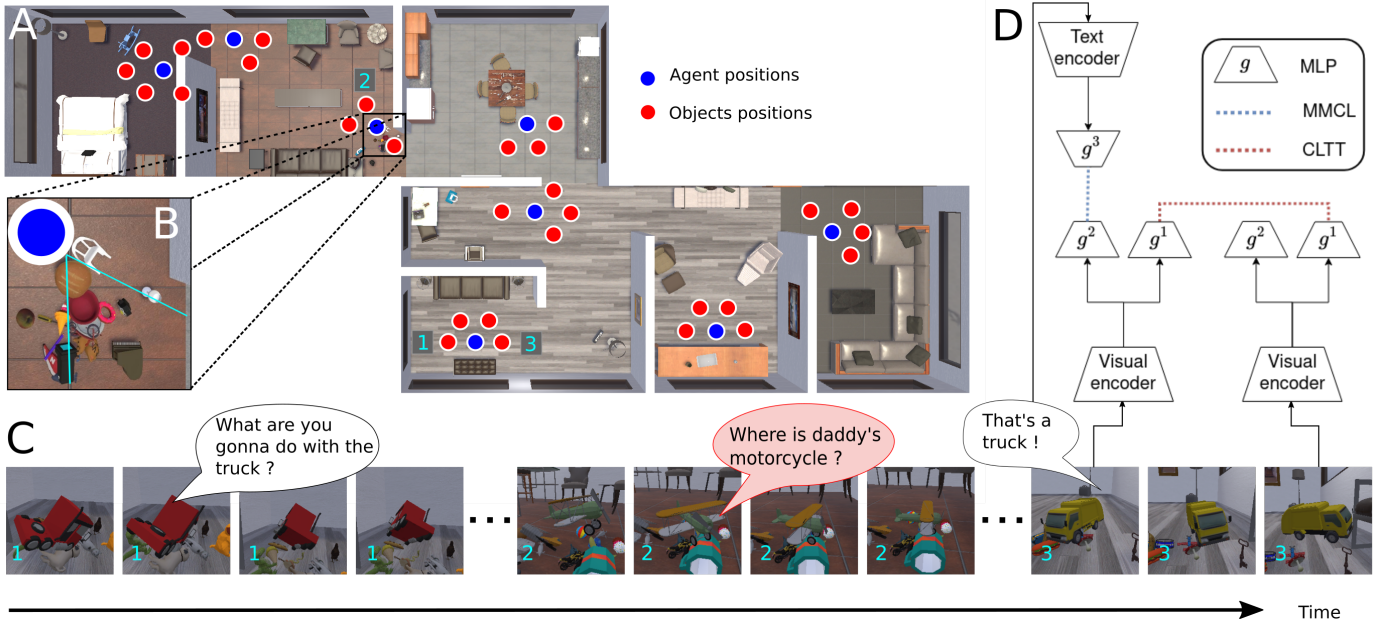


Fig. 1. A) Top view of the *Virtual Home Environment*, where blue and red dots, respectively, indicate possible agent and toy positions. The agent is always turned towards a toy position. Toy positions marked 1–3 correspond to sessions 1–3 in C. B) Zoom-in of the scene in A with turquoise lines indicating the agent’s field of view. C) Images extracted in a temporally ordered fashion (left to right) for three different “play” sessions of the Dyadic Play Dataset. Text boxes show examples of captions related to the manipulated object (white) or to another object in the background (red). D) Summary of the learning architecture, see Section III-B for details. Abbreviations: MLP: multi-layer perceptron, MMCL: multimodal contrastive learning, CLTT: contrastive learning through time.

II. RELATED WORK

a) Models of toddlers’ object learning: Recently, the availability of datasets of images extracted from toddlers’ head-mounted cameras [17], [21] has enabled the study of representations learnt through the senses of toddlers. This way, two recent models managed to learn word-vision mappings, but they either used a pretrained vision model [22], [23] or trained on curated data in a supervised fashion [23]. Another model extracted semantic visual representations by making similar close-in-time representations [24], but it did not leverage any sort of utterances. Working with such real-world ego-centric data has the limitation that a precise control of the statistics of the training data, like the sparsity or ambiguity of naming utterances, is difficult. This makes it hard to study how such statistics might affect toddlers’ object learning. A recent model also trained visual representations in cross-situational contexts, *i.e.* when naming utterances can refer to several visible objects [25]. However, they only consider synthetic images based on 9 digits rather than simulated ego-centric play sessions.

b) Slowness principle for learning visual representations: Several computational models showed that making similar close-in-time representations can lead to visual representation suitable for object instance recognition, scene recognition and object categorization [26], [27], [28], [29]. Integrating constraints from toddlers’ perceptual experience into these models, like short arms or foveation, can lead to improved ability to recognize objects in front of novel backgrounds [30]. None of these works modelled interactions with a caregiver whose utterances provide weak language supervision.

c) Visuo-language representation learning: Previous works explored multi-modal contrastive learning of audio-visual or text-visual representations and showed that language can supervise visual representation learning [31], [32], [33]. However, none of these works studied whether naming statistics in dyadic play sessions elicit semantic visual representations. A recent work showed that adding sparse labeling on top of contrastive learning through time enhances visual categorization [30]. Unlike them, we here consider textured objects in a cluttered home environment and developmentally-relevant statistics of caregivers’ utterances.

III. METHODS

In this section, we expose our computational model of visual representation learning during dyadic play. We first describe the Dyadic Play Dataset, which simulates dyadic play with objects (Section III-A). Then, we explain in Section III-B how we model a toddler’s learning process.

A. Dyadic Play Dataset (DPD)

The DPD contains 857,760 images of 224×224 pixels split into 42,888 “play sessions,” each containing 20 images. There are 12 sessions dedicated to each of 3574 objects [29], which are extracted from the Toys4k dataset [19]. To create them, we simulate a toddler that interacts with objects using the simulation platform ThreeDWorld (TDW) [34].

a) Image recording: At the beginning of each play session, we place the agent in a random location, with a random orientation within the *Virtual Home Environment* [29] (cf. Fig. 1A). We approximate an ego-centric view of a seated

toddler by positioning a camera at 0.4 m above the ground. At the beginning, the camera always watches the object being held, which we call “main object,” making the camera orientation dependent on the initial position of the main object. Because toddlers’ short arms constrain the way they hold objects, we randomly sample the relative starting position of the main object within a distance range of $[0.25; 0.35]$ m from the agent, $[-30; 30]^\circ$ on the sides and $[-30; -10]^\circ$ in elevation. Furthermore, toddlers tend to keep objects in an upright position [35], hence, we randomly orient the object around the yaw axis (unbounded) but bound the object orientation in $[-5; 5]^\circ$ around the pitch and roll axes.

In addition to the main object, we add 5 to 20 randomly sampled (without replacement) background objects to the scene, following the number of toys commonly present in the field of view of toddlers during play [17]. To place them, we follow the same procedure as for the main object, but change the range of distances from the agent to $[0.42; 0.875]$ m. In contrast to the object being held, we let them fall to the floor one at a time using the physics engine of TDW. Fig. 1B shows a top-view of a subsequent play area.

To simulate natural interactions, in each of the 20 frames of the play sessions the agent slowly moves and turns the object along/around the three axes. We display a high variability of manipulation sequences by generating all movements with an Ornstein-Uhlenbeck stochastic process for which we randomly generate a new recall coefficient in $[0.1, 1]$ at the beginning of each session, for each kind of manipulation and axis. Thus, the agent dynamically changes the manipulation speed within a session, as well as the amplitude of changes from one session to the next, independently for each axis of rotation/motion. We bound the motion speed by 3° per frame for side/elevation motion and 0.05 m for in-depth motion. However, the agent can not exceed the starting absolute bounds described above. Similarly, we bound the rotation speed by 20° around the yaw axis and 4° around the pitch and roll axis, also to simulate the upright bias of toddlers [35].

Finally, while the agent observes the moving object, it also executes additional relative small eye movements around the object in focus (yaw and pitch axis), whose endpoints are sampled from a Normal distribution $\mathcal{N}([0, 0], 2 \times \mathbb{I})$. We also allow the agent to focus its gaze on an object different from the main one with probability 0.7; however, we discard these images during training as previous methods already analyzed the impact of changes of attention on contrastive through time losses [27], [29]. Fig. 1D shows examples of the resulting sequences of toddler-centric views.

For evaluation purposes, we build a test dataset of 5 images per main object (a total of 17,870 images), each in different scenes generated as above. To avoid correlated images, we do not apply temporal manipulations used during training. Since rotations around pitch/roll axis may go beyond their initial boundaries after some manipulation time during training, it may create a subsequent data distribution shift between train and test images. Assuming that the pitch/roll rotations rarely

go beyond $[-20; 20]^\circ$ from their starting orientation during training, we increase the initial range of test object orientations to $[-20; 20]^\circ$ around the pitch and roll axis

b) *Naming utterances*: We aim to simulate developmentally-relevant statistics of utterances from a caregiver to a toddler during a play session. To achieve this, we source relevant transcripts from the CHILDES database [20]. Our study focuses specifically on transcripts from native English-speaking children residing in North America and the UK, between the ages of six months and two years. We only keep utterances of caregivers. As we are interested in object-related statements, we curate captions derived from statements identifying objects from the Toys4k dataset [19]. To form general templates, we only keep captions that occur across multiple object categories. After filtering out incomplete utterances and those containing less than three words, as well as a manual review, we retained 820 templates. Each of these templates can be used with any category name without relying on the context of the object.

B. Self-supervised learning of visual representations

In this work, we postulate that toddlers learn visual representations that 1) slowly change over time and 2) align with co-occurring linguistic representations. To model their learning process, we consider two previously introduced self-supervised loss functions. The first loss aligns embeddings of close-in-time images; this entails that seeing an object from different viewpoints elicits viewpoint invariant object representations [27]. The second loss function aligns embeddings of co-occurring images and naming utterances. Since the naming utterance may inform about the category of the object in focus, this provides weak supervision regarding the object category. To implement our learning mechanisms, we use SimCLR [36], a state-of-the-art contrastive learning algorithm. We sum up the learning architecture in Fig. 1E.

a) *Contrastive learning through time (CLTT)*: At each learning iteration, we sample a mini-batch \mathcal{X} that contains N randomly sampled images x_i . For each image x_i , we also randomly sample a single successor/predecessor x_j belonging to the same “play session” as x_i . Thus, x_i and x_j always display the same object, but observed through different orientations and positions. In addition, we expect that the active and multi-modal manipulation of the object biases toddlers’ attention onto the main object [12], [16]; thus, we apply on images, with probability 0.5, a center crop of size ranging in $[8; 100]\%$ of the image size. Thereafter, we compute embeddings of images $z^1 = g^1(f(x))$ using a feature extractor f and a projection head g^1 , both implemented as neural networks [36]. Finally, for a pair (z_i, z_j) , we minimize

$$l_T(z_i^1, z_j^1) = -\log \frac{e^{\cos(z_i^1, z_j^1)/\tau}}{\sum_{z \in Z^1, z \neq z_i^1} e^{\cos(z_i^1, z)/\tau}}, \quad (1)$$

where \cos stands for the cosine similarity, Z^1 contains all embeddings z^1 and τ is the temperature hyper-parameter [36]. The top part of (1) aims to move together image embeddings

that belong to the same manipulation session while the bottom part of (1) ensures that all image embeddings remain dissimilar.

b) Multimodal contrastive learning: For each sampled image x_i , we also sample, if provided by the caregiver, its co-occurring naming utterance l_i . We extract the pre-trained features of the naming utterances with a state-of-art text embedding model h [37]. Then, we compute different embeddings of images $z_i^2 = g^2(f(x_i))$ and captions $z_i^3 = g^3(h(l_i))$ using projection heads g^2, g^3 and a text feature extractor h in order to minimize, for each pair (z_i^2, z_i^3) :

$$l_M(z_i^2, z_i^3) - \log \frac{e^{\cos(z_i^2, z_i^3)/\tau}}{\sum_{\substack{z \in \mathcal{Z}^2 \\ z \neq z_i^2}} e^{\cos(z_i^2, z_k^2)/\tau}} \quad (2)$$

where \mathcal{Z}^2 contains the provided embeddings z^2 and z^3 . The top part of (2) ensures that co-occurring naming utterances and images have similar embeddings while the bottom part prevents all embeddings to collapse into a single vector.

The total loss function for a batch of size N is the symmetric and batch-wise sum of (1) and (2):

$$\frac{1}{2N} \sum_{i=1}^N l_M(z_i^3, z_i^2) + l_M(z_i^2, z_i^3) + l_T(z_i^1, z_j^1) + l_T(z_j^1, z_i^1).$$

C. Developmentally-relevant utterance statistics

To model developmentally-relevant utterance statistics, we extract statistics reported in at-home studies of toddlers' dyadic play sessions [13], [14]. First, [14] reports that 56% of naming events match the visual input of toddlers. Thus, we define the (conditional) probability of naming the manipulated object (versus another object in the background) as $p_{\text{correct}} = 0.5$. Second, a caregiver approximately provides, on average, one naming utterance related to the object being manipulated per time-extended object manipulation [13] (exactly $\frac{2.54}{2} = 1.27$ in their study). Since our play sessions last for 20 frames, we approximate the probability of naming the manipulated object during a frame as $p_{\text{name}} = \frac{1}{20} = 0.05$. We assume that consecutive frames are spaced by one second, making the duration of our play sessions correspond to the average of 20 seconds reported in [13]. Finally, it allows us to define the probability of naming any object in a frame $p_{\text{sparse}} = \frac{p_{\text{name}}}{p_{\text{correct}}} = 0.1$.

D. Training and evaluation

a) Training: We use a ResNet18 [38] as vision encoder f and, for all projection heads, a fully connected neural network with one hidden layer of size 256 followed by batch normalization and ReLU activation. For encoding the text, we use a pre-trained BERT [37] with 4 layers, 8 self-attention heads and a hidden size of 512, introduced as BERT-small by [39]. We train our models for 50 epochs with the AdamW optimizer, a learning rate of 0.001 and weight decay of 0.01. We tested a temperature hyper-parameter $\tau \in \{0.07, 0.1, 0.5\}$ and found 0.07 to be the best. The rationale behind using a pre-trained text model lies in our research focus on visual representation learning. Preliminary experiments (reported in

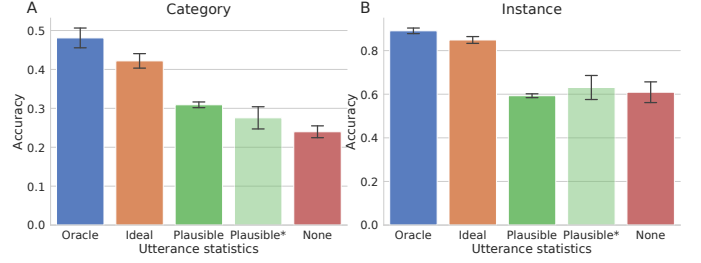


Fig. 2. A) Category recognition accuracy and B) object instance recognition accuracy for different settings. Oracle represents supervised learning, while an Ideal caregiver consistently names the correct object. Plausible stands for developmentally-relevant utterance statistics, Plausible* is identical but trains the text encoder from scratch.




Fig. 3. t-SNE visualization of the feature representations extracted by the vision- and text-encoder in different training settings. For better visualization, we show 10 random classes and a collection of randomly selected templates, each incorporating category-specific names.

experiments Section IV-A), demonstrate comparable recognition performance, although with a longer convergence time when training a randomly initialized text model concurrently with visual encoding.

b) Evaluation: To evaluate if the learned representation supports category recognition, we apply a repeated random subsampling cross-validation to split the non-overlapping object instances into 2382 train (2/3 of the total) and 1191 test objects (1/3 of the total). Then, we extract images of train and test objects to form the train and test datasets, respectively (see Section III-A). To evaluate our representation with respect to object instance recognition, we extract novel images of the train objects to build a test dataset. For both object instance and category recognition, we freeze the weights of the vision encoder and train a linear classifier in an online and supervised fashion on top of the latent visual representation [40].

IV. EXPERIMENTS & RESULTS

We first study whether developmentally-relevant utterance statistics support learning semantic visual representations. Then, we analyze the impact of the sparsity, ambiguity and the attention paid to the object's name on the learning process.

A. Developmentally relevant utterance statistics improve object representations

To assess the impact of plausible utterance statistics on the visual representation, Fig. 2 compares such statistics with two upper-bound baselines: Oracle refers to supervised learning and Ideal refers to a caregiver who always utters sentences




Fig. 4. Analysis of the impact of the sparsity parameters on A) category recognition and B) object instance recognition. The red points indicate the developmentally-relevant value.

containing the correct object. Our baseline (None) does not use any utterances relying solely on the time-contrastive loss for visual representations. We observe that our model improves category recognition over not using naming utterances, even though it does not reach the performance of oracles. Additionally, we examine how plausible utterance statistics perform when used to train the text representations jointly with the visual representations from scratch over 100 epochs. The model yields similar outcomes, albeit necessitating significantly more training time, as it has not yet reached convergence. We employ t-SNE visualization to validate our learned output feature representations. Fig. 3 shows that the models trained with utterances exhibit better separation than without textual guidance (baseline). We conclude that developmentally-relevant imperfect naming utterances can help to build semantic visual representation.

B. Small changes in utterance frequency and ambiguity drastically impact object representations

To understand how frequency and ambiguity of utterances affect object representations, we systematically vary both factors in Fig. 4 and Fig. 5, respectively. We observe that developmentally-relevant values (red points) are close to the inflection point of a sigmoidal function, where small shifts have a high impact on the quality of the learned representation. This suggests that toddler’s object learning may be quite sensitive to caregivers’ presence and the quality of their utterances. We also notice worse instance recognition with very few utterances in comparison to none. We suspect that these rare and ambiguous utterances mostly inject noise into the representation.

C. Attending to object names is crucial for learning good visual representations

To assess the importance of attention in acquiring semantic visual representations, we present in Fig. 6 a comparison of text embeddings across various utterances. Even with the infrequent application of the cross-modal loss, the model demonstrates the capability to group utterances sharing the same category name, clearly separating them from other utterance classes (as illustrated by the blocks in Fig. 6). Increasing the degree of textual guidance by small increments yields improved class separation. In addition, more good naming utterances amplify




Fig. 5. Analysis of the impact of naming ambiguity on A) category recognition and B) instance recognition. A high p_{correct} implies low ambiguity. The red points indicate the developmentally-relevant value.




Fig. 6. Analysis of the effect of the sparsity parameter and the correct naming probability p_{correct} on the textual representations. The heatmaps show the cosine similarity between the text embeddings of utterances from two random classes. The first element of each class is the raw category name.

linguistic attention on the name of the object being held, thereby allowing to better guide the visual representation. The distinct emerging clusters in Fig. 3 emphasise the models’ capacity to extract object names embedded within the utterances accurately. When considered alongside the accuracy values of classification experiments, this observation strongly suggests that increased textual attention on the object name is crucial for learning object representations.

V. CONCLUSION

We proposed a computational model of the development of toddlers object representations during dyadic play and investigated if and how caregivers’ utterances affect learned visual representations. We found that realistic utterance statistics elicit more semantic visual representations and the quality of these representations is sensitive to small shifts in the statistics of utterances. Our analysis also revealed that realistic statistics of utterances lead the model to focus on the object name within an utterance, suggesting that visual and language representations reciprocally affect each other.

The main limitation of our work is the gap between the actual experiences of toddlers and our simplified dataset. For instance, our “play sessions” are short and disconnected, as there is no notion of temporal proximity between different sessions and our captions represent a small subset of the space of possible utterances. In addition, hands, caregivers’ visual cues and the touch modality have been completely ignored. In the future, we plan to validate our findings on a dataset of real-world videos extracted from head-mounted cameras during dyadic play.

ACKNOWLEDGMENT

We gratefully acknowledge support from GENCI-IDRIS (Grant 2022-AD011014008) for providing computing and data-processing resources needed for this work. Additional support was received by the Deutsche Forschungsgemeinschaft (DFG project 5368 “Abstract REpresentations in Neural Architectures (ARENA)”), as well as the project “The Adaptive Mind” funded by the Excellence Program of the Hessian Ministry of Higher Education, Science, Research and Art (HMWK). JT was supported by the Johanna Quandt foundation.

REFERENCES

- [1] J. N. Wood and S. M. Wood, “The development of invariant object recognition requires visual experience with temporally smooth objects,” *Cognitive Science*, vol. 42, no. 4, pp. 1391–1406, 2018.
- [2] N. Li and J. J. DiCarlo, “Unsupervised natural visual experience rapidly reshapes size-invariant object representation in inferior temporal cortex,” *Neuron*, vol. 67, no. 6, pp. 1062–1075, 2010.
- [3] L. Wiskott and T. J. Sejnowski, “Slow feature analysis: Unsupervised learning of invariances,” *Neural computation*, vol. 14, no. 4, pp. 715–770, 2002.
- [4] K. S. Kraebel and P. C. Gerhardstein, “Three-month-old infants’ object recognition across changes in viewpoint using an operant learning procedure,” *Infant Behavior and Development*, vol. 29, no. 1, pp. 11–23, 2006.
- [5] K. C. Sloska and S. P. Johnson, “Development of three-dimensional object completion in infancy,” *Child development*, vol. 79, no. 5, pp. 1230–1236, 2008.
- [6] L. Gershkoff-Stowe and L. B. Smith, “Shape and the first hundred nouns,” *Child development*, vol. 75, no. 4, pp. 1098–1114, 2004.
- [7] L. B. Smith, S. S. Jones, B. Landau, L. Gershkoff-Stowe, and L. Samuelson, “Object name learning provides on-the-job training for attention,” *Psychological science*, vol. 13, no. 1, pp. 13–19, 2002.
- [8] S. S. Jones, “Late talkers show no shape bias in a novel name extension task,” *Developmental Science*, vol. 6, no. 6, pp. 477–483, 2003.
- [9] G. Diesendruck and P. Bloom, “How specific is the shape bias?” *Child development*, vol. 74, no. 1, pp. 168–178, 2003.
- [10] D. G. K. Nelson, A. Frankenfield, C. Morris, and E. Blair, “Young children’s use of functional information to categorize artifacts: Three factors that matter,” *Cognition*, vol. 77, no. 2, pp. 133–168, 2000.
- [11] A. Gopnik and A. Meltzoff, “The development of categorization in the second year and its relation to other cognitive and linguistic developments,” *Child development*, pp. 1523–1531, 1987.
- [12] S. E. Schroer and C. Yu, “Looking is not enough: Multimodal attention supports the real-time learning of new words,” *Developmental Science*, vol. 26, no. 2, p. e13290, 2023.
- [13] C. Suarez-Rivera, J. L. Schatz, O. Herzberg, and C. S. Tamis-LeMonda, “Joint engagement in the home environment is frequent, multimodal, timely, and structured,” *Infancy*, vol. 27, no. 2, pp. 232–254, 2022.
- [14] S. E. Schroer, R. E. Peters, A. Yarbrough, and C. Yu, “Visual attention and language exposure during everyday activities: an at-home study of early word learning using wearable eye trackers,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 44, no. 44, 2022.
- [15] E. M. Clerkin and L. B. Smith, “Real-world statistics at two timescales and a mechanism for infant learning of object names,” *Proceedings of the National Academy of Sciences*, vol. 119, no. 18, p. e2123239119, 2022.
- [16] J. Yang, L. Smith, D. Crandall, and C. Yu, “Using manual actions to create visual saliency: an outside-in solution to sustained attention and joint attention,” in *Proceedings of the Annual Meeting of the Cognitive Science Society*, 2023.
- [17] S. Bambach, D. J. Crandall, L. B. Smith, and C. Yu, “Active viewing in toddlers facilitates visual object learning: An egocentric vision approach,” in *CogSci*, 2016.
- [18] C. Yu, Y. Zhang, L. K. Slone, and L. B. Smith, “The infant’s view redefines the problem of referential uncertainty in early word learning,” *Proceedings of the National Academy of Sciences*, vol. 118, no. 52, p. e2107019118, 2021.
- [19] S. Stojanov, A. Thai, and J. M. Rehg, “Using shape to categorize: Low-shot learning with an explicit shape bias,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 1798–1808.
- [20] B. MacWhinney and C. Snow, “The child language data exchange system,” *Journal of child language*, vol. 12, no. 2, pp. 271–295, 1985.
- [21] J. Sullivan, M. Mei, A. Perfors, E. Wojcik, and M. C. Frank, “Saycam: A large, longitudinal audiovisual dataset recorded from the infant’s perspective,” *Open mind*, vol. 5, pp. 20–29, 2021.
- [22] W. K. Vong, E. Orhan, and B. Lake, “Cross-situational word learning from naturalistic headcam data,” in *34th CUNY Conference on Human Sentence Processing*, 2021.
- [23] S. Tsutsui, A. Chandrasekaran, M. A. Reza, D. Crandall, and C. Yu, “A computational model of early word learning from the infant’s point of view,” *arXiv preprint arXiv:2006.02802*, 2020.
- [24] E. Orhan, V. Gupta, and B. M. Lake, “Self-supervised learning through the eyes of a child,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 9960–9971.
- [25] W. K. Vong and B. M. Lake, “Cross-situational word learning with multimodal neural networks,” *Cognitive science*, vol. 46, no. 4, p. e13122, 2022.
- [26] M. Franzius, N. Wilbert, and L. Wiskott, “Invariant object recognition and pose estimation with slow feature analysis,” *Neural computation*, vol. 23, no. 9, pp. 2289–2323, 2011.
- [27] F. Schneider, X. Xu, M. R. Ernst, Z. Yu, and J. Triesch, “Contrastive learning through time,” in *SVRHM 2021 Workshop@ NeurIPS*, 2021.
- [28] C. Feichtenhofer, H. Fan, B. Xiong, R. Girshick, and K. He, “A large-scale study on unsupervised spatiotemporal representation learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 3299–3309.
- [29] A. Aubret, M. R. Ernst, C. Teulière, and J. Triesch, “Time to augment self-supervised visual representation learning,” in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=o8xdgmwCP8l>
- [30] A. Aubret, C. Teulière, and J. Triesch, “Toddler-inspired learning induces hierarchical object representations,” in *IEEE ICDL-Sensorimotor Interaction, language, and Embodiment of Symbols (SMILES) workshop*, 2022.
- [31] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al., “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [32] K. Desai and J. Johnson, “Virtex: Learning visual representations from textual annotations,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 11 162–11 173.
- [33] D. Harwath, A. Recasens, D. Surís, G. Chuang, A. Torralba, and J. Glass, “Jointly discovering visual objects and spoken words from raw sensory input,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 649–665.
- [34] C. Gan, J. Schwartz, S. Alter, D. Mrowca, M. Schrimpf, J. Traer, J. De Freitas, J. Kubilius, A. Bhandwaldar, N. Haber, et al., “Threed-world: A platform for interactive multi-modal physical simulation,” *arXiv preprint arXiv:2007.04954*, 2020.
- [35] A. F. Pereira, K. H. James, S. S. Jones, and L. B. Smith, “Early biases and developmental changes in self-generated object views,” *Journal of vision*, vol. 10, no. 11, pp. 22–22, 2010.
- [36] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, “A simple framework for contrastive learning of visual representations,” in *International conference on machine learning*. PMLR, 2020, pp. 1597–1607.
- [37] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 ACL*. Minneapolis, Minnesota: Association for Computational Linguistics, June 2019, pp. 4171–4186. [Online]. Available: <https://aclanthology.org/N19-1423>
- [38] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [39] I. Turc, M.-W. Chang, K. Lee, and K. Toutanova, “Well-read students learn better: On the importance of pre-training compact models,” 2019.
- [40] F. Bordes, R. Balestrierio, and P. Vincent, “Towards democratizing joint-embedding self-supervised learning,” *arXiv preprint arXiv:2303.01986*, 2023.