



**HAL**  
open science

# Toddler-inspired embodied vision for learning object representations

Arthur Aubret, Céline Teulière, Jochen Triesch

► **To cite this version:**

Arthur Aubret, Céline Teulière, Jochen Triesch. Toddler-inspired embodied vision for learning object representations. IEEE ICDL, Sep 2022, Londres, United Kingdom. hal-03838291

**HAL Id: hal-03838291**

**<https://hal.science/hal-03838291>**

Submitted on 3 Nov 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Toddler-inspired embodied vision for learning object representations

Arthur Aubret  
University Clermont Auvergne  
CNRS, Pascal institute  
Clermont-Ferrand, France  
arthur.aubret@uca.fr

Céline Teulière  
University Clermont Auvergne  
CNRS, Pascal institute  
Clermont-Ferrand, France  
celine.teuliere@uca.fr

Jochen Triesch  
Frankfurt Institute for Advanced Studies  
Frankfurt am Main, Germany  
triesch@fias.uni-frankfurt.de

**Abstract**—Recent time-contrastive learning approaches manage to learn invariant object representations without supervision. This is achieved by mapping successive views of an object onto close-by internal representations. When considering this learning approach as a model of the development of human object recognition, it is important to consider what visual input a toddler would typically observe while interacting with objects. First, human vision is highly foveated, with high resolution only available in the central region of the field of view. Second, objects may be seen against a blurry background due to toddlers’ limited depth of field. Third, during object manipulation a toddler mostly observes close objects filling a large part of the field of view due to their rather short arms. Here, we study how these effects impact the quality of visual representations learnt through time-contrastive learning. To this end, we let a visually embodied agent “play” with objects in different locations of a near photo-realistic flat. During each play session the agent views an object in multiple orientations before turning its body to view another object. The resulting sequence of views feeds a time-contrastive learning algorithm. Our results show that visual statistics mimicking those of a toddler improve object recognition accuracy in both familiar and novel environments. We argue that this effect is caused by the reduction of features extracted in the background, a neural network bias for large features in the image and a greater similarity between novel and familiar background regions. The results of our model suggest that several influences on toddler’s visual input statistics support their unsupervised learning of object representations.

**Index Terms**—embodied vision, contrastive learning, unsupervised learning, object recognition, foveation, depth of field

## I. INTRODUCTION

Babies learn to build representations of their surroundings before knowing the word for an object, suggesting that they perform a form of unsupervised learning [1]. What mechanism underlies their learning process? Several experiments show that modifying the temporal contiguity of an animal’s visual experiences shapes its visual representations such that temporally close visual inputs tend to be similarly represented in the visual cortex [2], [3]. For example, by slowly rotating an object, one learns to extract the information shared between the views, *i.e.*, the object identity [3]. The overall idea that the brain uses temporal closeness to build visual representations without supervision has been called the *slowness principle* [4], [5]. One way to implement the slowness principle in machine learning is through *contrastive learning*. While the original

idea of contrastive learning was to learn visual representations (almost) invariant to data-augmentations like image crops or rotations, one can also frame data-augmentations as a shift in time [6]; in fact, time-based augmentations like object rotations [6] or saccades [7] allow to build task-useful representations with performance close to less natural augmentations [7]. Therefore, contrastive learning with time-based augmentations [6] appears to be a strong candidate for understanding how animals construct invariant object representations.

In current machine learning research, contrastive learning methods mostly consider ImageNet-like datasets of high-resolution images with a great variety of differently-sized objects [8]. Yet, because of its embodiment, a toddler experiences very different kinds of visual inputs. First, the density of cones is higher in the fovea compared to the periphery, leading to poor resolution in peripheral vision [9]. Second, a child learns to control its distance of focus within a few months [10]. Due to the limited depth of field, a properly focused object will be seen against a more or less blurry background. Third, the ability of a toddler to act onto its environment and the environment itself shape what the toddler sees [11]. A typical toddler experiences a relatively small set of rooms/outdoor environments [12], spending most of its time in the home or a daycare environment. Furthermore, 18-months toddlers have short arms, making objects they play with appear bigger in their field of view in comparison with adults [13], [14].

Here, we assume that the brain applies the slowness principle to learn internal representations without supervision. Under this assumption, we aim to study whether and how the above-mentioned visual statistics impact the learnt visual representations. To control different aspects of the visual statistics and investigate their impact, we visually embody a virtual agent using a near photo-realistic simulation platform, ThreeDWorld (TDW) [15] and a recent dataset of 3D models of toy objects [16]. We introduce a house environment in which we simulate this agent playing with different objects in different locations: we simulate object rotations, saccades, and switching between objects. These transformations induce natural time-based augmentations for training with time-contrastive learning. Our first experimental contribution demonstrates that the diversity/complexity of backgrounds can greatly hurt object recognition, thereby motivating the



Fig. 1. Scene of TDW used to situate our agent in a house. In blue, we indicate possible agent positions and we display in red possible object positions. In this episode, the agent is located in the office and can interact with four objects.




Fig. 2. Submarine with different orientations, perceived through different eye orientation, but in front of a single background. From left to right, we illustrate: a low amplitude of saccades that keeps the object centered; the progressive blur of the scene according to its relative depth with the focused object (depth of field); the progressive blur as pixels get distant from the image center (foveation); the increase in size that follows the greatest closeness of the object; the combination of all transformations.

need for attenuating background information. In our second experimental contribution, we investigate the impact of four particular visual statistics on object recognition through time-contrastive learning: the foveation blurring the surroundings of a fixated object; the depth of field effect blurring objects in the background; the distance of the object from the agent; the range of saccades during object exploration. Our analysis is three-fold: 1) we show that an object-centered foveation and a restricted depth of field promote object representations that generalize to novel environments. However, they do **not** improve object representations in familiar environments; 2) we exhibit that making small saccades attenuates the extraction of background information; 3) presumably due to a neural network bias, closer objects make easier object recognition, in line with previous results [13]. The combination of all three effects drastically improves object recognition accuracy, thereby showing that toddler-inspired visual inputs improve object representations.

## II. RELATED WORKS

*a) SimCLR-TT:* SimCLR, a state-of-the-art unsupervised contrastive learning algorithm [8], recently managed to learn representations of visual images useful to downstream classification tasks. Intuitively, two semantically similar/dissimilar images should be respectively close/distant in the latent space, such that the distance between images in the latent space captures the semantics of the images. Following the slowness principle, SimCLR-TT [6], [17] expects two semantically similar images to be successive in time. SimCLR-TT samples an image  $x_i$  and selects its successive or previous image (in time) to be its positive sample  $x_j$ . Then, SimCLR-TT samples

$N - 1$  other images  $x_k$  different from  $x_i$ , and their predecessor/successor as  $x_{N+k}$ . Thereafter, SimCLR-TT minimizes the loss function of SimCLR by adapting the parameters  $\theta$  of the representation encoder (a neural network)  $f_\theta$ . Thus, it minimizes the following loss function for all positive pairs:

$$\mathcal{L}_{i,j} = -\log \frac{e^{\text{sim}(f_\theta(x_i), f_\theta(x_j))/\tau}}{\sum_{k=1, k \neq i, k \neq j}^{2N} e^{\text{sim}(f_\theta(x_i), f_\theta(x_k))/\tau}}, \quad (1)$$

where  $\text{sim}$  and  $\tau$  are respectively a similarity function and the temperature hyper-parameter. The numerator of eq. 1 makes sure that successive images are close in the latent space. To avoid the collapse of the representation into a single vector, the denominator makes sure that all representations are relatively distant.

*b) Visual data-augmentations:* Most contrastive learning approaches use unrealistic data augmentations like color distortion, crop, rotation, resize, blur [8], Jigsaw [18], selection of color channels [19]. In practice, it has been argued that these augmentations are related to a subset of more realistic augmentations [20]: crop can relate to occlusion, resize to object depth motion, color changes to illumination changes etc. Whether unrealistic augmentations and their realistic counterpart truly result in similar representations remains unclear. [21] proposed as a data-augmentation to detect the salient object and switch its background with one from another object; the resulting improved accuracy validates that discarding background information may improve the robustness of object representations. Another body of work tries to learn embeddings of videos thanks to the temporal contiguity of frames. Then, positive pairs are different clips (temporal crops) of the same video [22]–[24] or overlapping clips with different temporal windows [25]. We refer to [26] for a review of contrastive learning methods along with the different data-augmentation methods. Unlike previous visual contrastive learning approaches, we visually embody our agent in a 3D environment that allows for both ego-centric and object-centric transformations.

*c) Time-contrastive learning for object representations:* closer to our work, we emphasize three works that study the representations learnt through realistic transformations. Using videos, some approaches proposed to take advantage of the temporal information to learn image embeddings [27], [28]. But datasets may not simulate the depth of field/foveation and do not allow the study of the complexity of the background. Similarly to us, [6] consider 3D object rotations as transformations, but with a uniform background and without ego-motions. Interestingly, they show that long-time fixations improve the representations in comparison with short-time fixations. [7] showed that foveation based on magnification and saccades can efficiently replace crop augmentations. Unlike us, they do not study the impact of visual statistics on the learnt representation.

*d) Toddler-inspired machine learning:* while we are not aware of works that studied our set of visual statistics with toddlers, several works already studied the impact of some visual statistics on machine learning methods with respect to

different objectives. However, they do not consider the context of time-contrastive learning. [13] showed that supervised learning could benefit from learning with images of objects, sized and centered as experienced by toddlers; in fact, agents may even move to create un-occluded views of an object to improve recognition [29]. Machine learning simulations also support that the increase of acuity in the first months of life provides a natural curriculum for learning. This can aid the learning of reflexes, like the acquisition of vergence behavior (alignment of the two eyes on the same point) [30] or the detection of binocular disparity (difference between the two eye images) [31]. This can also aid learning representations of faces robust to different resolutions [32], [33].

### III. METHOD

We want to simulate the visual observations experienced by toddlers playing with objects in order to study the impact of visual statistics on representations learnt through SimCLR-TT [6]. To simulate the visual environment of toddlers, we consider a house environment, as illustrated in Figure 1, and manually define different play locations. These locations can have similar/different floors, walls and objects, thereby making non-obvious their recognition. In a location, the agent engages in 100-step long play sessions (also called episodes), which are initially built in a 3-steps process. Firstly, we place the agent in the center of a randomly sampled location. Secondly, we put randomly sampled objects around the agent (between two and six according to the location and surrounding objects). Because of their different positions, a different background gets associated to each of them. The objects are randomly sampled from a set of 20 (unless stated otherwise) textured and untextured objects randomly extracted from the Toys-4k dataset. We expect untextured objects to make the task harder, since the model can not focus on the color regardless of the shape of the object. Finally, we randomly turn the agent in front of an object. For example, in Figure 1, the agent runs a play session in the office of the house and can play with four objects (each seen against a different background).

To simulate viewing sequences experienced by a child during a play session, we introduce three factors of variation that change across the time of a play session. At a low frequency (every 10 timesteps), the agent selects a new object (the first one to the right of the current object) and rotates its body in the direction of the object. At a higher frequency (every timestep), it rotates the object by an azimuth angle uniformly sampled in  $[0; 30]$  and executes a saccadic eye movement by sampling its position from a normal distribution  $S \times \mathcal{N}([0, 0], \mathbb{I})$  where  $S$  is a hyper-parameter denoting the saccade amplitude. Note that, at  $[0, 0]$ , the eyes are aligned with the object such that it appears centered in the field of view. We expect this difference of frequency to make the agent discard the information about high-frequency variations (saccades and object rotation) while keeping information about low-frequency variations [6], *i.e.*, object identity, background identity and location identity. The highly frequent transformations have different properties that we expect to differently impact the learnt features: the object

rotation minimally impacts the background and the ego-centric saccades impacts both the object and the background.

We now introduce the three visual effects shaping the agent’s visual input, namely foveation, distance/centering of objects, and depth of field.

#### A. Foveation and saccades

In humans, photo-receptors called cones are distributed on the retina of each eye to detect light under day-light conditions. These cones are present with high density in the fovea and low density in the periphery of the retina [34] such that humans have high visual accuracy in the center of gaze and low accuracy in the periphery (**foveation**). To simulate this change of acuity, we use a previous model of humans’ visual acuity [35]. It creates a pyramid of multi-resolution images and interpolates each pixel between the different resolutions according to the distance of each pixel from the focus position.

To simulate different amplitudes of **saccades**, and thus different levels of overlap between the object and the high-accuracy area, we consider a range saccade amplitudes  $S \in \{3, 2, 1\}$ . For  $S = 3$ , this results in a range of saccade angles (yaw and pitch axis) with significant density in  $[-15^\circ, 15^\circ]$  degrees from the object center. Low amplitudes of saccades result in a strong average overlap between the high-accuracy area and the object position, and inversely for high amplitudes of saccades. In practice, human saccades amplitudes distribution seems to be close to an inverse Gaussian distribution [36], however, we expect our simple distribution of saccades to be qualitatively consistent with empirical findings showing many small and few large saccades.

#### B. Depth of field

The accommodation reflex of humans describes their ability to focus on an object of interest depending on its distance from the observer in order to obtain a sharp image of the object. An important part of the reflex is the increase of the curvature of each eye lens, which adapts the refractive effect of the lens on the light to the desired focus distance. For instance, if one starts to look at the sky right after reading a book, the lens will decrease its curvature to have the distance of focus matching the distance of the sky. Out of focus objects will get progressively blurred as they lie on a plane distant from the focus plane. The range of depths around the distance of focus observed with *acceptably sharp focus* is called the depth of field. In practice, according to light levels or the depth of the object of focus, the iris increases/decreases the depth of field by respectively decreasing/increasing the size of the pupil. To keep it simple, we set the focus distance of the camera to correspond to the distance of the manipulated toy and we decrease the depth of field by decreasing the **aperture** number of the TDW platform [15]. Overall, this makes the objects appear in focus while blurring the background (cf. Figure 2).

#### C. Object position/distance

Between one and two years, toddlers start to walk, manipulate and play with objects. This learning step considerably

changes what a child sees; while it was previously seeing very cluttered scenes of relatively small objects, it starts to watch close and centered objects [13], [14]. To simulate this change of object distance, we bring objects **closer** to the agent by a factor of 0.7. Figure 2 illustrates the size change that results from decreasing the distance between the object and the agent. We do not assume this distance to be very realistic, but a true hand-based manipulation of the object is out of our scope and our choice avoids physical collisions. However, this modification allows us to qualitatively assess the impact of modifying the agent-to-object distance.

#### IV. EXPERIMENTS

We aim to evaluate whether considering the visual statistics of toddlers improves object recognition in time-contrastive learning, considering the typical environment of a toddler.

During training, the agent accumulates observations in a buffer of size 80,000 and, between each timestep, it trains the encoder network with SimCLR-TT [6]. The  $128 \times 128$  RGB images go through a succession of convolution layers with the following [channels, kernel size, stride, padding] structure: [64, 8, 4, 2], [128, 4, 2, 1], [256, 2, 2, 1], [256, 2, 2, 1]. These are followed by an average pooling layer and a linear layer ending with 20 units (unless stated otherwise). After each convolution layer, we apply a non-linear ReLU and a dropout layer ( $p = 0.5$ ) to prevent over-fitting. Our preliminary results showed that such a shallow architecture combined with Euclidean distance, dropout and average pooling, works better than standard architectures used in contrastive learning [8]. We hypothesize that this stems from the use of less diverse inputs. We apply a weight decay of  $1e - 6$  and update weights with the AdamW optimizer [37] and a learning rate of  $5e - 4$ .

For evaluation purposes, we train linear classifiers on top of the learnt representations to assess their quality [8]. We either train them to recognize the objects or the backgrounds. We use two datasets of at least 1200 images (depending on the number of backgrounds): 1) a validation dataset composed of images generated similarly to the training data; 2) a test dataset composed of similar object views in unseen backgrounds (from another house). We apply the visual effects (foveation, depth of field, etc.) on the training images and on the validation and test datasets. In all experiments, we compute the mean and standard deviation over five random seeds.

##### A. Diverse complex backgrounds impede object learning

To study the effects of different backgrounds on object learning, we set up one location with six object positions/backgrounds and we progressively increase the diversity/complexity of features in the six backgrounds. Specifically, we train our representation with: fully white background almost without clues about what background is being watched; an untextured empty room with floor/wall corners apparent; the same empty room with an oriented parquet; a room with the same parquet, textured walls and several objects. The left part of Figure 4 suggests that increasing the complexity of the background has a minor effect on object representations,




Fig. 3. Saliency maps of the SimCLR loss on a bus toy seen against a white background (left) or a room with several background objects (right)

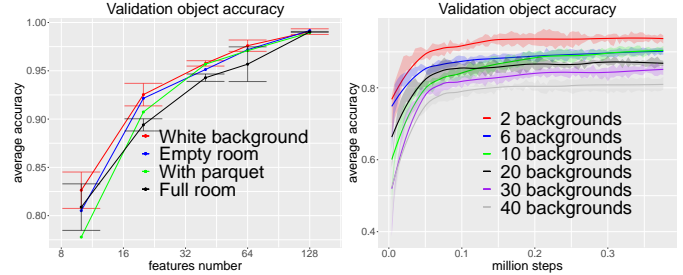


Fig. 4. **Left:** Object recognition accuracy on the validation dataset as a function of dimensionality of the latent space for different levels of background complexity in the single-room environment, after 400k training steps/interactions. Error bars indicate the standard deviation of the *white background* and *full room* experiments, we omit those of other curves for clarity. **Right:** Object recognition accuracy on the validation dataset for different numbers of backgrounds seen by the agent. Shaded regions around the curves indicate one standard deviation.

mostly for a medium size latent space (size of 20 and 40 and 64 neurons).

To investigate the impact of the number of backgrounds on the representation, we set up different houses with increasingly higher number of different backgrounds (the sum of backgrounds for all possible locations). The right part of Figure 4 highlights that, as we increase the overall number of backgrounds, the object recognition accuracy decreases. To visualize what the encoder is extracting from the image depending on the number/complexity of the background, we display in Figure 3 typical saliency map [38] outputs by the encoder when applying the SimCLR loss, depending on the background of the image. We clearly observe that the encoder better focuses on the object when there is no distracting background. Overall, these experiments suggest that object recognition accuracy can be improved by reducing distracting information from the background.

##### B. Toddler-inspired visual input aids object learning

We train the representation of our agent with/without the visual effects of §III on the environment displayed in Figure 1 and evaluate the representation on both validation and test datasets. In Figure 5, we apply one transformation at a time. We first note that the effect of the visual statistics remain consistent over different sizes of latent space.

a) *Closer objects:* Figure 5 (top) shows a steady increase of object recognition accuracy in both validation/test datasets with closer objects. We observe an opposite decrease in the background accuracy (bottom). Importantly, we observe in Figure 6 (left) better object representations (test dataset) even without different backgrounds; thus parts of the accuracy



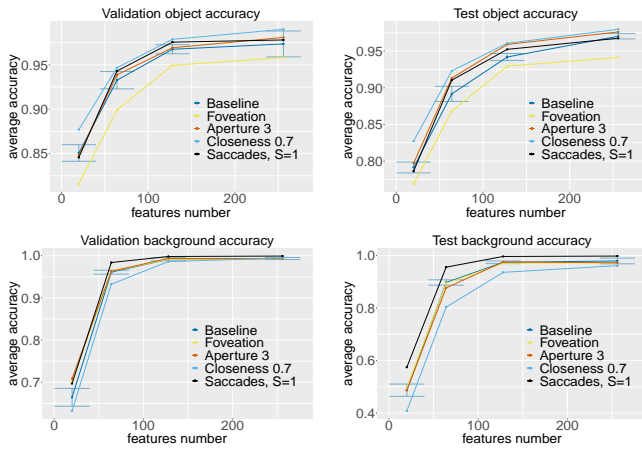


Fig. 5. **Top:** Object recognition accuracy according to the applied visual statistics on the validation dataset (left) and test dataset (right). **Bottom:** Background recognition accuracy according to the applied visual statistics on the validation dataset (left) and test dataset (right). The baseline corresponds to perfectly clear images (No foveation and aperture of 20) of relatively distant objects. The horizontal bars indicate the standard deviation of the baseline, we omit those of other curves for clarity.

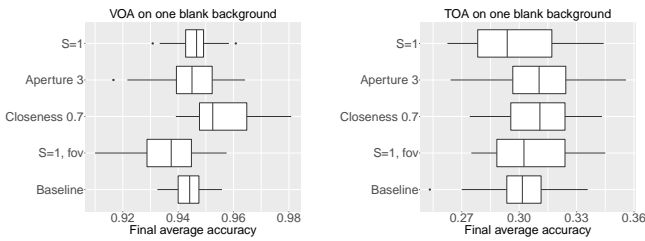


Fig. 6. Final average object accuracy on the validation dataset (VOA, left) and the test dataset (TOA, right) with a representation trained on one blank background. We took the last 10 average accuracies (each 4000 steps between 360 000 and 400 000 timesteps) of the five seeds to compute the statistics. Since there is only one background, the agent does not turn its body but selects a new object every 10 timesteps to keep the same frequency of object switches as in other experiments.

increase does not depend on background information. We deduce that part of the improvement comes from our CNN, presumably biased to better recognize larger objects. This is consistent with previous works [13] on supervised learning.

*b) Low Saccade amplitudes:* Figure 5 shows: 1) on the top-left graph, an increase of object recognition accuracy in both validation/test datasets with medium size of latent spaces (64 and 128); 2) on bottom graphs, a steady increase in the background recognition accuracy in validation/test datasets. As small saccades do not impact the representation when there is no background (see Figure 6), this increased object recognition accuracy does not come from the reduced number of training/test diversity of examples (due to the lower amplitudes of spatial shifts). Instead, saccades generate higher shifts in backgrounds than on the objects, due to the higher distance of the background. Consequently, with high-amplitude saccades, a *chair* visible in a background may be seen in one view, but not in another one. An agent needs to extract more than

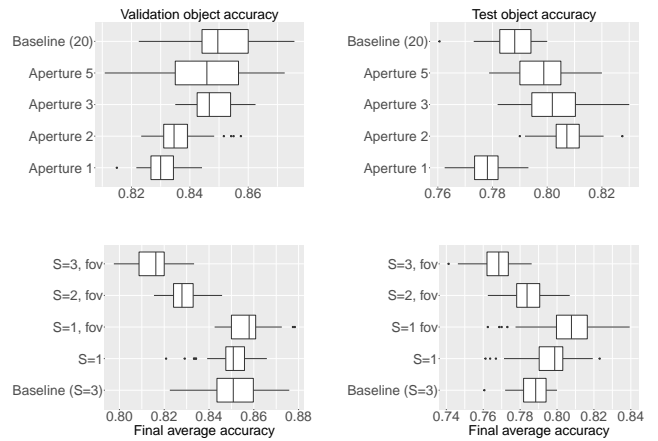


Fig. 7. Object recognition accuracy on the validation dataset (left) and test dataset (right). We took the last 10 average accuracies (each 4000 steps between 360 000 and 400 000 timesteps) of the five seeds to compute the statistics. **Top:** Accuracy according to the number of aperture. Aperture 20 refers to the baseline. **Bottom:** Accuracy according to saccade amplitudes and the foveation of images.



Fig. 8. Illustration of different images of airplanes extracted from our validation dataset, with different saccade positions and backgrounds. From left to right, we use an aperture of 1, 2, 3, 5, 20.

the *chair* feature from the background to learn the similarity between views. Therefore, we conclude that a low amplitude of saccades reduces the number of features needed to learn saccade-invariant background representations, allowing the extraction of more object-oriented features. In compliance with our analysis in §IV-A, this effect vanishes with too-small or too-large latent spaces.

*c) Limited depth of field:* Figure 5 (top-right) exhibits a sound increase of object recognition accuracy with an aperture of 3 only in the test datasets. As this increase does not appear in the validation set, we deduce that it does not come from the attenuation of background information. To investigate its origin, we compare the impact of different aperture numbers on the object recognition accuracy in Figure 7 (top). We observe a monotonic improvement on the validation object accuracy as we increase the aperture, saturating with an aperture of 3 and an important decrease with a low aperture (Aperture 1). This suggests that a too-low aperture, which blurs the object, prevents the extraction of important object features. We also observe a sweet spot for an aperture of 2 or 3 on the test object accuracy. The highest performances are thus obtained with a blurred background and slightly blurred parts of the object (see Figure 8). To check whether the improvement comes from blurring small parts of the object or blurring the background, we trained the representation in the house environment using completely unblurred objects (Aperture 5 in Figure 7, top-




Fig. 9. Test object recognition accuracy using toddler-inspired visual inputs on 20 (left) and 80 (right) objects. The *All* experiment combines all image transformations. We set the aperture to 4 instead of 3 in the *All* experiments to account for the increase of depth of field induced by closer objects in humans.

right) and in front of a single white background (Aperture 3 in Figure 6, right). A clear object in a blurred background already shows an improved test object recognition accuracy, and a blurred object in a clear background do not significantly impact the representation. Thus, we conclude that blurring the background is crucial to obtain this improved generalization. We hypothesize that a depth-wise blur increases the similarity of backgrounds between test and validation datasets, thereby increasing the generalization of objects on new backgrounds.

*d) Foveated images:* Figure 5 (top) shows that foveation can lead to worse object representation. To study the origin of this decrease, in Figure 7 (bottom), we vary the range of amplitudes of saccades to control how much saccades direct the gaze towards the object or the background. A low amplitude makes, on average, saccades closer to the object and inversely. We observe that the foveation increases the object recognition accuracy only when the saccade amplitudes are small, such that most saccades end up on the object. Fixation of the object is thus crucial when considering foveated inputs. Similarly to the decrease of the depth of field, object-focused foveation improves object recognition accuracy on the test dataset, but not on the validation set. As for the depth of field, we expect it comes from more similar features between test/validation backgrounds. This tends to be confirmed in Figure 9 (left), since the addition of aperture upon foveated images does not significantly improve the object recognition accuracy.

*e) Combining all visual statistics:* in Figure 9, we progressively apply all visual statistics and assess the quality of the object representations. We vary the number of objects to check whether our results scale with an increased object number. Overall, we clearly see that using toddler-inspired visual inputs improves object representations learnt through time-contrastive learning.

## V. CONCLUSION

We hypothesized that time-contrastive learning, as an implementation of the slowness principle, models the unsupervised learning of visual representations in toddlers. We simulated viewing sequences as might be experienced by a toddler by letting an agent move and interact with different objects in different locations of a house. First, we showed that the complexity and the number of backgrounds, negatively

impact object representations learnt through time-contrastive learning in toddler-like environments. This suggests that object representations are improved by attenuating the information from the background. We then showed that incorporating viewing effects of a toddler aids the learning of object-oriented representations in a house environment. Our analysis suggests that foveation and a restricted depth of field have a similar effect of making object representations more robust to novel environments. We did not observe a consequent attenuation of background information. We also discovered that object-centered saccades are crucial to avoid the deterioration of the object representations due to foveation. Closer objects improve object recognition, presumably because they induce a greater bias towards object-centered features in neural networks [13]. However, while we did not show that closer objects attenuate the background information, we did not show either that this attenuation does not exist. Finally, using time-contrastive learning, a low amplitude of saccades seems to limit the extraction of features from the background. We showed that these three effects are cumulative. Overall, our contribution highlights the need to rethink the usual framework of contrastive learning as an embodied learning process.

As we found that low-amplitude saccades are the only visual condition that indeed slightly attenuates the extraction of background information, it remains a large margin of improvement in learning background-invariant representations. The solution may be to consider different frequencies of objects/backgrounds switches, since time-contrastive learning is expected to favour encoding slowly varying features. Furthermore, the incorporation of covert attention mechanisms seems promising.

Our study focuses on a small set of factors of variation in small environments, but other aspects also impact a toddler’s visual input. For example, we do not consider other agents in the environment, the navigation of the agent between different rooms, cluttered scenes or the visibility of arms/hands of the agent. Furthermore, we assume the agent executes normally distributed saccades and rotates objects independently from their view, which are inconsistent with findings from the developmental literature [36], [39]. Playing is rather an active learning process where the toddler watches and manipulates what he or she is interested in, depending on their actual visual inputs. As such, integrating intrinsic motivations [40] to generate actions in an unsupervised way may be crucial for modeling the development of toddlers’ object recognition skills.

## ACKNOWLEDGMENT

This research was supported by “The Adaptive Mind” and “The Third Wave of Artificial Intelligence”, funded by the Excellence Program of the Hessian Ministry of Higher Education, Science, Research and Art. This work was sponsored by a public grant overseen by the French National Agency through the IMobS3 Laboratory of Excellence (ANR-10-LABX-0016) and the IDEX-ISITE initiative CAP 20-25 (ANR-16-IDEX-0001). Financial support was also received from Clermont Au-

vergne Metropole through a French Tech-Clermont Auvergne professorship. We gratefully acknowledge support from the CNRS/IN2P3 Computing Center (Lyon - France) for providing computing and data-processing resources needed for this work. This work was also performed using HPC resources from GENCI-IDRIS (Grant 2022-AD011011623R1). JT is supported by the Johanna Quandt Foundation. We thank Markus Ernst for his help with the foveation code.

## REFERENCES

- [1] P. C. Quinn, P. D. Eimas, and S. L. Rosenkrantz, "Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants," *Perception*, vol. 22, no. 4, pp. 463–475, 1993.
- [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] 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.
- [4] L. Wiskott and T. J. Sejnowski, "Slow feature analysis: Unsupervised learning of invariances," *Neural computation*, vol. 14, no. 4, pp. 715–770, 2002.
- [5] A. Hyvarinen and H. Morioka, "Unsupervised feature extraction by time-contrastive learning and nonlinear ica," *Advances in neural information processing systems*, vol. 29, 2016.
- [6] F. Schneider, X. Xu, M. R. Ernst, Z. Yu, and J. Triesch, "Contrastive learning through time," in *SVRHM 2021 Workshop@ NeurIPS*, 2021.
- [7] B. Wang, D. Mayo, A. Deza, A. Barbu, and C. Conwell, "On the use of cortical magnification and saccades as biological proxies for data augmentation," in *SVRHM 2021 Workshop@ NeurIPS*, 2021.
- [8] 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.
- [9] H. Strasburger, I. Rentschler, and M. Jüttner, "Peripheral vision and pattern recognition: A review," *Journal of vision*, vol. 11, no. 5, pp. 13–13, 2011.
- [10] D. C. Currie and R. E. MANNY, "The development of accommodation," *Vision Research*, vol. 37, no. 11, pp. 1525–1533, 1997.
- [11] L. Byrge, O. Sporns, and L. B. Smith, "Developmental process emerges from extended brain–body–behavior networks," *Trends in cognitive sciences*, vol. 18, no. 8, pp. 395–403, 2014.
- [12] 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*, pp. 1–10, 2020.
- [13] S. Bambach, D. Crandall, L. Smith, and C. Yu, "Toddler-inspired visual object learning," *Advances in neural information processing systems*, vol. 31, 2018.
- [14] L. B. Smith, C. Yu, and A. F. Pereira, "Not your mother's view: The dynamics of toddler visual experience," *Developmental science*, vol. 14, no. 1, pp. 9–17, 2011.
- [15] C. Gan, J. Schwartz, S. Alter, D. Mrowca, M. Schrimpf, J. Traer, J. De Freitas, J. Kubilius, A. Bhandwadar, N. Haber *et al.*, "Threedworld: A platform for interactive multi-modal physical simulation," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
- [16] 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*, 2021, pp. 1798–1808.
- [17] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [18] I. Misra and L. v. d. Maaten, "Self-supervised learning of pretext-invariant representations," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6707–6717.
- [19] Y. Tian, C. Sun, B. Poole, D. Krishnan, C. Schmid, and P. Isola, "What makes for good views for contrastive learning?" *Advances in Neural Information Processing Systems*, vol. 33, pp. 6827–6839, 2020.
- [20] A. Laflaquière, "A sensorimotor perspective on contrastive multiview visual representation learning," *IEEE Transactions on Cognitive and Developmental Systems*, 2021.
- [21] C. Ryali, D. J. Schwab, and A. S. Morcos, "Learning background invariance improves generalization and robustness in self-supervised learning on imagenet and beyond," 2021.
- [22] 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.
- [23] R. Qian, T. Meng, B. Gong, M.-H. Yang, H. Wang, S. Belongie, and Y. Cui, "Spatiotemporal contrastive video representation learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 6964–6974.
- [24] J. Wang, G. Bertasius, D. Tran, and L. Torresani, "Long-short temporal contrastive learning of video transformers," *arXiv preprint arXiv:2106.09212*, 2021.
- [25] A. Recasens, P. Luc, J.-B. Alayrac, L. Wang, F. Strub, C. Tallec, M. Malinowski, V. Pătrăucean, F. Althé, M. Valko *et al.*, "Broaden your views for self-supervised video learning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1255–1265.
- [26] A. Jaiswal, A. R. Babu, M. Z. Zadeh, D. Banerjee, and F. Makedon, "A survey on contrastive self-supervised learning," *Technologies*, vol. 9, no. 1, p. 2, 2021.
- [27] J. Knights, B. Harwood, D. Ward, A. Vanderkop, O. Mackenzie-Ross, and P. Moghadam, "Temporally coherent embeddings for self-supervised video representation learning," in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 8914–8921.
- [28] E. Orhan, V. Gupta, and B. M. Lake, "Self-supervised learning through the eyes of a child," *Advances in Neural Information Processing Systems*, vol. 33, pp. 9960–9971, 2020.
- [29] J. Yang, Z. Ren, M. Xu, X. Chen, D. J. Crandall, D. Parikh, and D. Batra, "Embodied amodal recognition: Learning to move to perceive objects," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2040–2050.
- [30] A. Priamikov, V. Narayan, B. E. Shi, and J. Triesch, "The role of contrast sensitivity in the development of binocular vision: A computational study," in *2015 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*. IEEE, 2015, pp. 33–38.
- [31] M. Dominguez and R. A. Jacobs, "Developmental constraints aid the acquisition of binocular disparity sensitivities," *Neural Computation*, vol. 15, no. 1, pp. 161–182, 2003.
- [32] H. Jang and F. Tong, "Convolutional neural networks trained with a developmental sequence of blurry to clear images reveal core differences between face and object processing," *Journal of vision*, vol. 21, no. 12, pp. 6–6, 2021.
- [33] L. Vogelsang, S. Gilad-Gutnick, E. Ehrenberg, A. Yonas, S. Diamond, R. Held, and P. Sinha, "Potential downside of high initial visual acuity," *Proceedings of the National Academy of Sciences*, vol. 115, no. 44, pp. 11 333–11 338, 2018.
- [34] Á. Szél, P. Röhlich, A. R. Caffé, and T. Van Veen, "Distribution of cone photoreceptors in the mammalian retina," *Microscopy research and technique*, vol. 35, no. 6, pp. 445–462, 1996.
- [35] M. Jiang, S. Huang, J. Duan, and Q. Zhao, "Salicon: Saliency in context," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1072–1080.
- [36] H. A. Trukenbrod, S. Barthelmé, F. A. Wichmann, and R. Engbert, "Spatial statistics for gaze patterns in scene viewing: Effects of repeated viewing," *Journal of Vision*, vol. 19, no. 6, pp. 5–5, 2019.
- [37] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2018.
- [38] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *In Workshop at International Conference on Learning Representations*. Citeseer, 2014.
- [39] 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.
- [40] A. Aubret, L. Matignon, and S. Hassas, "A survey on intrinsic motivation in reinforcement learning," *arXiv preprint arXiv:1908.06976*, 2019.