
AVoE: A Synthetic 3D Dataset on Understanding Violation of Expectation for Artificial Cognition

Arijit Dasgupta¹, Jiafei Duan², Marcelo H. Ang Jr¹, and Cheston Tan²

¹Department of Mechanical Engineering, National University of Singapore

²Institute for Infocomm Research, A*STAR

Abstract

Recent work in cognitive reasoning and computer vision has engendered an increasing popularity for the Violation-of-Expectation (VoE) paradigm in synthetic datasets. Inspired by work in infant psychology, researchers have started evaluating a model’s ability to discriminate between expected and surprising scenes as a sign of its reasoning ability. Existing VoE-based 3D datasets in physical reasoning only provide vision data. However, current cognitive models of physical reasoning by psychologists reveal infants create high-level abstract representations of objects and interactions. Capitalizing on this knowledge, we propose AVoE: a synthetic 3D VoE-based dataset that presents stimuli from multiple novel sub-categories for five event categories of physical reasoning. Compared to existing work, AVoE is armed with ground-truth labels of abstract features and rules augmented to vision data, paving the way for high-level symbolic predictions in physical reasoning tasks. AVoE: <https://github.com/jiafei1224/AVoE>

1 Introduction

Physical reasoning systems built on the foundations of intuitive physics & psychology are pivotal to creating machines that learn and think like humans [1, 2]. The ability to physically reason like humans forms a crucial gateway to multiple real-world applications ranging from robotic assistants, autonomous vehicles, and safe AI tools. These systems are guided by facets of core knowledge [3], from object representation to numerical, spatial & social reasoning. Infant psychologists theorised that human newborns have an in-built physical reasoning [4, 5] and object representation [6, 7, 8] system guiding their learning of core knowledge. This has inspired researchers to approach the creation of physical reasoning systems as a start-up software embedded with core knowledge principles [9].

Predicting the future of a physical interaction given a scene prior has been the most common task in designing computational physical reasoning systems. One such task is the tower of falling objects, which has been as a test-bed for evaluating intuitive physics engines [10, 11, 12]. There have also been recent advancements in creating benchmarked datasets and models that combine multiple physical prediction tasks in a 3D environment [13, 14, 15]. Physical reasoning has also been explored in Embodied AI [16] & Visual Question Answering (VQA) [17] with public datasets like **CLEVRER** [18], **CRAFT** [19] & **TIWIQ** [20] providing scenes of general and random interactions between objects and multiple questions concerning the physical outcome.

A parallel track complementary to future prediction tasks is the design of artificial agents that can measure the plausibility of physical scenes. An agent capable of physical reasoning should not only be able to predict the future but also recognise if a scene is *possible* or *impossible*. The Violation-of-Expectation (VoE) paradigm is an empirical diagnostic tool first implemented in developmental psychology studies [21, 22] to measure the surprise of infants when shown *possible* or *impossible* scenes. The studies found that infants as young as 4 months could express surprise at a constructed

scene that violated the principle of object permanence: the tendency for an object to continually exist without dropping in and out of existence. This was akin to a magic show, and the surprise was measured by the looking time of the infants. VoE has since been used in an array of infant psychology experiments on a range of physical reasoning event categories [23, 24, 25, 26, 27].

The work in VoE has encouraged recent computational development of models and datasets [28, 29, 30] that challenge artificial agents to differentiate between *possible* and *impossible* scenes in physical events. Instead of training with *possible* and *impossible* scenes, only *possible* scenes & interactions are shown in the training dataset consistently among the studies. This inherently replicates the experience of a normal growing infant who is usually exposed to *possible* scenes. Therefore, the added challenge of an imbalanced and realistic training dataset means that researchers have to look beyond general end-to-end deep learning models and focus on generative and probabilistic approaches leveraging on inductive biases. VoE is not limited to physical reasoning, for it can be used in any reasoning task that has a clear binary distinction between *surprising* and *expected*. **AGENT** [31] & **BIB** [32] are two similar and recently developed datasets that train and evaluate an agent’s ability to recognise another agent’s goal preference [33], action efficiency [34] and more. Taking inspiration from the increasing use of VoE for computational reasoning tasks, we propose AVoE: a novel VoE-based 3D dataset of five different event categories of physical reasoning. Our contributions for AVoE are two-fold:

- Existing VoE-based physical reasoning datasets [28, 29, 30] only provide vision-based data. While vision still forms the main sensing modality of AVoE, we take a step further by providing additional scene-wise ground-truth metadata of abstract features and rules which models can exploit for higher-level symbolic prediction of the physical reasoning tasks.
- Compared to the existing datasets, AVoE contains scene stimuli in multiple novel sub-categories for each event category, adding to its diversity and complexity.

2 Related Work

The intersection of computer vision and physical reasoning is heavily grounded in the literature of psychology. The age in which infants adopt core principles of persistence, inertia & gravity [35] have been widely studied in a variety of event categories. For example, the **barrier** event is an event category illustrating (or violating) the constraint of solidity. In studies that implemented the barrier event to infants [23, 36], psychologists placed a solid barrier with an object on one side and the *surprising* event occurred when the infant was made to believe that the object could pass through the barrier. Like the barrier event, there are other events like **containment** [37, 26, 38] where a *surprising* scene was demonstrated with a taller object being fully contained in a shorter container, violating solidity constraints. Continuity constraints were also violated in **occluder** events [39, 24] when objects teleported behind one occluder to the back of another disconnected occluder. This event was also modified with different occluder heights in the middle segment [40]. Researchers also experimented with inertia violations in **collision** events [27, 41] where the outcome of the contact was impossible via linear momentum. Finally, gravity violations were also studied as **support** events [42, 25, 5] where imbalanced objects that appeared stable were presented to infants.

We find that **Piloto et al.** [28], **IntPhys** [29] & **ADEPT** [30] to be the most relevant to our work, as they all employ the VoE paradigm in their datasets and evaluation metrics. Additionally, they all present 3D datasets of very similar event categories of physical reasoning.

Piloto et al. [28] presents a 3D VoE dataset of 100,000 training videos and 10,000 pair probes of *surprising* and *expected* videos for evaluation. The dataset categorized their videos into ‘object persistence’, ‘unchangeableness’, ‘continuity’, ‘solidity’ & ‘containment’. Their Variational Autoencoder model benchmarked on the dataset and showed promise in ‘assimilating basic physical concepts’.

IntPhys [29] is a 3D VoE dataset with 15,000 videos of *possible* events and 3,960 videos of *possible & impossible* events in the test and dev sets. Only three events on ‘object permanence’, ‘shape constancy’ & ‘continuity’ were present. The study benchmarked the performances of a convolutional autoencoder and generative adversarial network with short & long-term predictions. The models performed poorly in comparison with their adult human trials but with higher than chance performance.

ADEPT [30] is a model that uses deep recognition networks, extended probabilistic simulation and particle filtering to predict expectations of objects. They use this model on a 3D VoE dataset of 1,000 training videos of random objects colliding and 1,512 test videos of *surprising* or *control* stimuli.

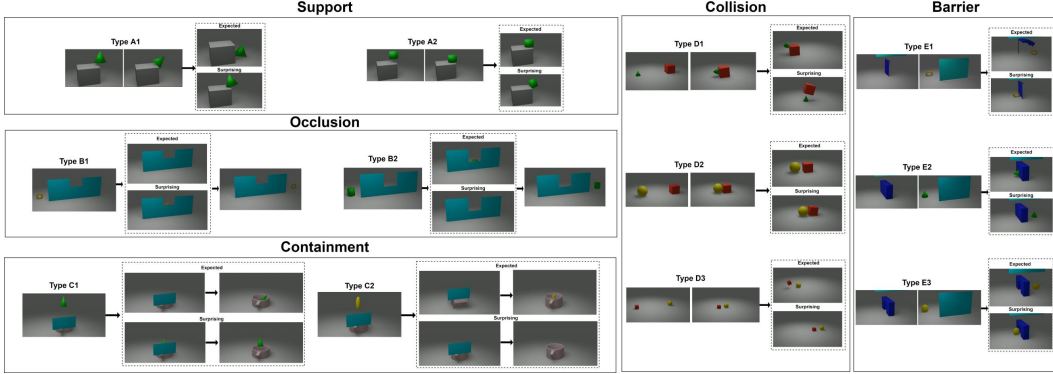


Figure 1: Examples of the different event categories in AVoE with their *surprising* and *expected* outcomes. **Support:** Type A1 (*originally unbalanced*), Type A2 (*originally balanced*). **Occlusion:** Type B1 (*object shorter than occluder*), Type B2 (*object taller than occluder*). **Containment:** Type C1 (*Object fully contained in container*), Type C2 (*object protruding out of container*). **Collision:** Type D1 (*same speed, different size*), Type D2 (*same speed, same size*), Type D3 (*different speed, same size*). **Barrier:** Type E1 (*soft barrier*), Type E2 (*solid barrier*), Type E3 (*barrier with opening*).

ADEPT accurately predicts the expected location of objects behind an occluder to measure surprise, while replicating adult human judgements on the ‘how, when & what’ traits of *surprising* scenes [43].

A common finding among developmental psychologists on physical reasoning is that infants create abstract representations of objects [6, 7] from which they extract spatial features. These high-level features are coupled with high-level rules of reasoning that infants develop over time via process known as explanation-based learning [44] to form their expectation on how a physical scene should play out [35]. For instance, the features of a containment event could refer to the heights of the object and container. A rule can be presented as a question: “is the object taller than the container?” to which the answer ‘yes’ (based on a simple comparison of the features) would signal an outcome of the object protruding out of the container and vice versa. This guided the construction of AVoE, which provides the ground-truth values for these rules and features in every scene.

3 AVoE Dataset

3.1 Overview

Figure 1 comprehensively illustrates the composition of the AVoE dataset, which comprises synthetic video sub-datasets in five event categories: **support** (A), **occlusion** (B), **containment** (C), **collision** (D) & **barrier** (E). Each one of these event categories are split into further sub-categories (e.g. Type B2 etc.) that showcase physical variations based on the differing stimuli of the scene. For each subcategory, every trial showcases an *expected* or *surprising* version of the stimuli. All frames were developed in the open-source 3D graphics software Blender [45], using a Python API. Every event category types $\psi \in \{A, B, C, D, E\}$ comes with sets of abstract features f^ψ , prior rules r_{prior}^ψ & posterior rules r_{post}^ψ . Prior rules are physical conditions about the event which can be answered with the features (e.g. height, width). These prior rules suffice to answer posterior rules that represent the outcome of the physical interaction. f^ψ , r_{prior}^ψ & r_{post}^ψ are only used for training as VoE evaluation does not test a model’s prediction of them. Still, they are still provided for the test and validation sets should researchers wish to evaluate performance in predicting features, prior & posterior rules.

Support (Type A): An object is dropped right above the edge of the object. The object’s centre of mass is situated either over the edge (**Type A1**) or within the edge (**Type A2**).

Occlusion (Type B): An inert object has an initial momentum to pass behind an occluder. The object can either be shorter than the occluder’s middle portion (**Type B1**) or taller (**Type B2**).

Containment (Type C): An inert object falls from a short height above a container, with the interaction hidden behind an occluder. The object is short enough to be fully contained inside the container (**Type C1**) or tall enough to protrude out of the container top (**Type C2**).




Figure 2: Outcome distribution for AVoE for each one of the event categories (**D** has 2 outcomes). The distribution depicts ‘yes’ or ‘no’ answers to a question about each outcome for every trial.

Collision (Type D): Two inert objects with initial momentum collide head-on. In the first case (**Type D1**), two objects have the same initial speed with different sizes. In the other two cases, both objects have similar size, with either the different (**Type D2**) or same (**Type D3**) initial speeds.

Barrier (Type E): An inert object has an initial momentum to pass through a barrier, with their interaction hidden behind an occluder. To experiment with different barriers, AVoE comprises events with either a soft barrier (**Type E1**), a solid barrier (**Type E2**) or a barrier with opening (**Type E3**).

3.2 Procedural Generation

Multiple physical stimuli that affect the outcome of the interaction were randomly sampled to amplify the diversity of the dataset. Common parameters among the 5 sub-datasets included the object’s shape $S_{Obj} \in \{Cube, Cylinder, Torus, Sphere, Cone, Side\ Cylinder, Inverted\ Cone\}$ and the object’s height & width, $H_{Obj}, W_{Obj} \in [0.4, 1.6]$, where a 1 is equivalent to $2m$ in the 3D environment. The initial contact point of the object on the support in **A** $C_{Obj} \in [0.2, 0.8]$ where $C_{Obj} = 0.2$ indicates that the 20% of the object’s width is over the edge. In **B**, the occluder’s middle segment height, $H_{Occ} \in [0.1, 0.9]$ with 1 being the height of the occluder. In **C**, the container’s shape, $S_{Con} \in \{Mug, Box\}$ was also varied with its height and width, $H_{Con}, W_{Con} \in [0.5, 1.5]$. The height ($\propto mass^{\frac{1}{3}}$) and initial speeds of objects in **D** were sampled as $H_{Obj} \in [0.5, 1.5]$ & $V_{Obj} \in [0.5, 2.5]$. In **E3**, the barrier’s opening height and width were sampled in the range, $H_{Bar}, W_{Bar} \in [0.4, 1.4]$.

3.3 Dataset Structure

Each event category ψ has 500 different configured trials, amounting to a total of 2,500 configurations in AVoE. 75% of the trials form the training set which comprises only *expected* videos. 15% and 10% of the trials are dedicated to the validation and test sets respectively. The validation and test sets contain both *expected* and *surprising* video pairs with the same stimuli in each trial, keeping the balance at an even 50%. This sums to 3125 videos. The appendix specifies the raw distribution of AVoE. At 50 frames per video, AVoE offers 156,250 frames, each having a size of 960×540 . Besides the RGB videos, AVoE also provides the depth map and instance segmented frames. Along with the automatically generated ground-truth labels of f^ψ, r_{prior}^ψ & r_{post}^ψ in every video, the frame-wise world position and orientation of all entities are provided. Figure 2 shows the outcome distribution for each of the 500 stimuli per event category, highlighting that the outcomes are relatively balanced. These outcomes are not forcefully set, rather, they result from the diverse variation of the stimuli.

4 Conclusion

In this work, we showcase an alternative approach to predicting VoE in basic physical reasoning events. By leveraging on findings in psychology literature, we propose AVoE, a novel synthetic 3D dataset augmented with ground-truth labels of abstract features and rules for five basic event categories of physical reasoning. These abstract features and rules can be structurally embedded into models as inductive biases to train an artificial agent’s ability to discern VoE in physical scenes. AVoE also further splits each event category stimuli into sub-categories for added diversity and complexity. Our recommended evaluation metrics for AVoE are the relative error rate and absolute error rate ($1 - AUC$) as described in [29]. As future work, we plan to benchmark the performances of baseline generative and probabilistic models along with a novel approach exploiting the feature and rule ground truths of AVoE.

Acknowledgements

This research is supported by the Agency for Science, Technology and Research (A*STAR), Singapore under its AME Programmatic Funding Scheme (Award #A18A2b0046). We would also like to sincerely thank Distinguished Professor Renée L. Baillargeon, Dr Yi Lin and Professor Su-Hua Wang for their consultation and expertise in the domain of infant psychology. Their advice has allowed for the integration of core infant psychology concepts with AI and the creation and improvement of AVoE.

References

- [1] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, “Building machines that learn and think like people,” *Behavioral and brain sciences*, vol. 40, 2017.
- [2] S. Adams, I. Arel, J. Bach, R. Coop, R. Furlan, B. Goertzel, J. S. Hall, A. Samsonovich, M. Scheutz, M. Schlesinger *et al.*, “Mapping the landscape of human-level artificial general intelligence,” *AI magazine*, vol. 33, no. 1, pp. 25–42, 2012.
- [3] E. S. Spelke and K. D. Kinzler, “Core knowledge,” *Developmental science*, vol. 10, no. 1, pp. 89–96, 2007.
- [4] R. Baillargeon, J. Li, Y. Gertner, and D. Wu, “How do infants reason about physical events?” 2011.
- [5] S. J. Hespos and R. Baillargeon, “Young infants’ actions reveal their developing knowledge of support variables: Converging evidence for violation-of-expectation findings,” *Cognition*, vol. 107, no. 1, pp. 304–316, 2008.
- [6] D. Kahneman, A. Treisman, and B. J. Gibbs, “The reviewing of object files: Object-specific integration of information,” *Cognitive psychology*, vol. 24, no. 2, pp. 175–219, 1992.
- [7] R. D. Gordon and D. E. Irwin, “What’s in an object file? evidence from priming studies,” *Perception & Psychophysics*, vol. 58, no. 8, pp. 1260–1277, 1996.
- [8] M. Stavans, Y. Lin, D. Wu, and R. Baillargeon, “Catastrophic individuation failures in infancy: A new model and predictions,” *Psychological Review*, vol. 126, no. 2, p. 196–225, 2019.
- [9] T. D. Ullman and J. B. Tenenbaum, “Bayesian models of conceptual development: Learning as building models of the world,” *Annual Review of Developmental Psychology*, vol. 2, pp. 533–558, 2020.
- [10] P. W. Battaglia, J. B. Hamrick, and J. B. Tenenbaum, “Simulation as an engine of physical scene understanding,” *Proceedings of the National Academy of Sciences*, vol. 110, no. 45, pp. 18 327–18 332, 2013.
- [11] R. Zhang, J. Wu, C. Zhang, W. T. Freeman, and J. B. Tenenbaum, “A comparative evaluation of approximate probabilistic simulation and deep neural networks as accounts of human physical scene understanding,” *arXiv preprint arXiv:1605.01138*, 2016.
- [12] A. Lerer, S. Gross, and R. Fergus, “Learning physical intuition of block towers by example,” in *International conference on machine learning*. PMLR, 2016, pp. 430–438.
- [13] J. Duan, S. Y. B. Jian, and C. Tan, “Space: A simulator for physical interactions and causal learning in 3d environments,” *arXiv preprint arXiv:2108.06180*, 2021.
- [14] D. M. Bear, E. Wang, D. Mrowca, F. J. Binder, H.-Y. F. Tung, R. Pramod, C. Holdaway, S. Tao, K. Smith, L. Fei-Fei *et al.*, “Physion: Evaluating physical prediction from vision in humans and machines,” *arXiv preprint arXiv:2106.08261*, 2021.
- [15] J. Duan, S. Yu, S. Poria, B. Wen, and C. Tan, “Pip: Physical interaction prediction via mental imagery with span selection,” *arXiv preprint arXiv:2109.04683*, 2021.
- [16] J. Duan, S. Yu, H. L. Tan, H. Zhu, and C. Tan, “A survey of embodied ai: From simulators to research tasks,” *arXiv preprint arXiv:2103.04918*, 2021.
- [17] Q. Wu, D. Teney, P. Wang, C. Shen, A. Dick, and A. van den Hengel, “Visual question answering: A survey of methods and datasets,” *Computer Vision and Image Understanding*, vol. 163, pp. 21–40, 2017.

- [18] K. Yi, C. Gan, Y. Li, P. Kohli, J. Wu, A. Torralba, and J. B. Tenenbaum, “Clevrer: Collision events for video representation and reasoning,” *arXiv preprint arXiv:1910.01442*, 2019.
- [19] T. Ates, M. S. Atesoglu, C. Yigit, I. Kesen, M. Kobas, E. Erdem, A. Erdem, T. Goksun, and D. Yuret, “Craft: A benchmark for causal reasoning about forces and interactions,” *arXiv preprint arXiv:2012.04293*, 2020.
- [20] M. Wagner, H. Basevi, R. Shetty, W. Li, M. Malinowski, M. Fritz, and A. Leonardis, “Answering visual what-if questions: From actions to predicted scene descriptions,” in *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018, pp. 0–0.
- [21] R. Baillargeon, E. S. Spelke, and S. Wasserman, “Object permanence in five-month-old infants,” *Cognition*, vol. 20, no. 3, pp. 191–208, 1985.
- [22] R. Baillargeon, “Object permanence in 31/2- and 41/2-month-old infants,” *Developmental psychology*, vol. 23, no. 5, p. 655, 1987.
- [23] R. Baillargeon, M. Graber, J. Devos, and J. Black, “Why do young infants fail to search for hidden objects?” *Cognition*, vol. 36, no. 3, pp. 255–284, 1990.
- [24] E. S. Spelke, R. Kestenbaum, D. J. Simons, and D. Wein, “Spatiotemporal continuity, smoothness of motion and object identity in infancy,” *British journal of developmental psychology*, vol. 13, no. 2, pp. 113–142, 1995.
- [25] N. Dan, T. Omori, and Y. Tomiyasu, “Development of infants’ intuitions about support relations: Sensitivity to stability,” *Developmental Science*, vol. 3, no. 2, pp. 171–180, 2000.
- [26] S.-h. Wang, R. Baillargeon, and S. Paterson, “Detecting continuity violations in infancy: A new account and new evidence from covering and tube events,” *Cognition*, vol. 95, no. 2, pp. 129–173, 2005.
- [27] L. Kotovsky and R. Baillargeon, “Calibration-based reasoning about collision events in 11-month-old infants,” *Cognition*, vol. 51, no. 2, pp. 107–129, 1994.
- [28] L. Piloto, A. Weinstein, D. TB, A. Ahuja, M. Mirza, G. Wayne, D. Amos, C.-c. Hung, and M. Botvinick, “Probing physics knowledge using tools from developmental psychology,” *arXiv preprint arXiv:1804.01128*, 2018.
- [29] R. Riochet, M. Y. Castro, M. Bernard, A. Lerer, R. Fergus, V. Izard, and E. Dupoux, “Int-phys: A framework and benchmark for visual intuitive physics reasoning,” *arXiv preprint arXiv:1803.07616*, 2018.
- [30] K. Smith, L. Mei, S. Yao, J. Wu, E. Spelke, J. Tenenbaum, and T. Ullman, “Modeling expectation violation in intuitive physics with coarse probabilistic object representations,” *Advances in Neural Information Processing Systems*, vol. 32, pp. 8985–8995, 2019.
- [31] T. Shu, A. Bhandwaldar, C. Gan, K. Smith, S. Liu, D. Gutfreund, E. Spelke, J. Tenenbaum, and T. Ullman, “Agent: A benchmark for core psychological reasoning,” in *Proceedings of the 38th International Conference on Machine Learning*, vol. 139. PMLR, 2021, pp. 9614–9625.
- [32] K. Gandhi, G. Stojnic, B. M. Lake, and M. R. Dillon, “Baby intuitions benchmark (bib): Discerning the goals, preferences, and actions of others,” *arXiv preprint arXiv:2102.11938*, 2021.
- [33] A. L. Woodward, “Infants selectively encode the goal object of an actor’s reach,” *Cognition*, vol. 69, no. 1, pp. 1–34, 1998.
- [34] G. Gergely, Z. Nádasdy, G. Csibra, and S. Bíró, “Taking the intentional stance at 12 months of age,” *Cognition*, vol. 56, no. 2, pp. 165–193, 1995.
- [35] Y. Lin, M. Stavans, and R. Baillargeon, “Infants’ physical reasoning and the cognitive architecture that supports it,” 2020.
- [36] E. S. Spelke, K. Breinlinger, J. Macomber, and K. Jacobson, “Origins of knowledge,” *Psychological review*, vol. 99, no. 4, p. 605, 1992.
- [37] S. J. Hespos and R. Baillargeon, “Infants’ knowledge about occlusion and containment events: A surprising discrepancy,” *Psychological Science*, vol. 12, no. 2, pp. 141–147, 2001.
- [38] Y. Mou and Y. Luo, “Is it a container? young infants’ understanding of containment events,” *Infancy*, vol. 22, no. 2, pp. 256–270, 2017.

- [39] R. Baillargeon and J. DeVos, “Object permanence in young infants: Further evidence,” *Child development*, vol. 62, no. 6, pp. 1227–1246, 1991.
- [40] A. Aguiar and R. Baillargeon, “2.5-month-old infants’ reasoning about when objects should and should not be occluded,” *Cognitive psychology*, vol. 39, no. 2, pp. 116–157, 1999.
- [41] L. Kotovsky and R. Baillargeon, “The development of calibration-based reasoning about collision events in young infants,” *Cognition*, vol. 67, no. 3, pp. 311–351, 1998.
- [42] R. Baillargeon, A. Needham, and J. DeVos, “The development of young infants’ intuitions about support,” *Early development and parenting*, vol. 1, no. 2, pp. 69–78, 1992.
- [43] K. Smith, L. Mei, S. Yao, J. Wu, E. S. Spelke, J. Tenenbaum, and T. D. Ullman, “The fine structure of surprise in intuitive physics: when, why, and how much?” in *CogSci*, 2020.
- [44] R. Baillargeon and G. F. DeJong, “Explanation-based learning in infancy,” *Psychonomic bulletin & review*, vol. 24, no. 5, pp. 1511–1526, 2017.
- [45] B. O. Community, *Blender - a 3D modelling and rendering package*, Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. [Online]. Available: <http://www.blender.org>

Appendix

Table 1: The number of trials for each event category in AVoE

Set	Support (A)			Occlusion (B)			Containment (C)			Collision (D)				Barrier (E)				Total
	A1	A2	Total	B1	B2	Total	C1	C2	Total	D1	D2	D3	Total	E1	E2	E3	Total	
Train	180	195	375	159	216	375	168	207	375	170	34	171	375	73	63	239	375	1875
Val	35	40	75	35	40	75	48	27	75	37	4	34	75	14	12	49	75	375
Test	26	24	50	24	26	50	31	19	50	16	9	25	50	8	15	27	50	250
Total	241	259	500	218	282	500	247	253	500	223	47	230	500	95	90	315	500	2500