PNAS RESEARCH ARTICLE | NEUROSCIENCE PSYCHOLOGICAL AND COGNITIVE SCIENCES

The world through infant eyes: Evidence for the early emergence of the cardinal orientation bias

Zachary J. Petroff[a,b,c], Swapnaa Jayaraman[d,e], Linda B. Smith[d], T. Rowan Candy[a], and Kathryn Bonnen[a,1]

Edited by J. Anthony Movshon, New York University, New York, NY; received November 12, 2024; accepted February 22, 2025

The structure of the environment includes more horizontal and vertical (i.e. cardinal) orientations than oblique orientations, meaning that edges tend to be aligned with or perpendicular to the direction of gravity. This bias in the visual scene is associated with a bias in visual sensitivity in adults. Although infants must learn to function in this biased environment, their immature motor control prevents them from consistently orienting themselves, relative to gravity. This study therefore asked whether cardinal orientations dominate human visual experience from early infancy or only from later in development, as motor control improves. We analyze video clips from head-mounted cameras, showing the egocentric perspective of 75 infants (1 to 12 mo) in their home environments in two communities (Indiana, USA vs. Tamil Nadu, India). We measured the distribution of orientations in each frame of these videos and found that horizontal and vertical orientations were overrepresented in infants from both countries. A cardinal orientation bias was evident even in the egocentric view of the youngest infants (3 wk) and became more prominent during the subsequent weeks of development. The early presence of a cardinal orientation bias in infants' visual input may serve as a consistent cue to gravity and ground planes, potentially influencing motor development and contributing to the formation of sensory, perceptual, and cognitive biases.

egocentric | natural scene statistics | development | infant vision | cardinal orientation bias

Natural events on Earth are influenced by gravity, establishing a reference frame with two primary axes: a vertical axis aligned to the direction of gravity and a horizontal axis defined by the ground plane, perpendicular to gravity. These two primary axes play a role in many observed phenomena. For example, they impact the phototropic growth of plants (1), the stability of objects, and the posture and movement of biological organisms (2–4). As a result, vertical and horizontal structures are overrepresented in natural and man-made visual environments (5). This overrepresentation of horizontal and vertical information is present at similar or even greater levels in typical "pastoral" natural scenes compared to indoor and outdoor "carpentered" scenes with man-made components (5). This overrepresentation is also reflected in human visual perception and cognition, including in performance on perceptual tasks (6–10) and in the neurophysiological function of mature primate visual systems (11–14).

Humans exhibit a well-documented "oblique effect," performing better on perceptual tasks (e.g., detection, discrimination) at horizontal and vertical orientations than at oblique orientations (6). These performance asymmetries reflect biased sensitivities that align with the statistics of natural and man-made scenes (5, 7, 15, 16). In other words, humans seem to be more sensitive to the orientations that they see most frequently. Strong arguments have been presented that these biases are the result of evolutionary pressures on vision in a world governed by gravity (7, 17).

Just as perceptual sensitivities align with the statistics of visual scenes, developmental studies indicate that both neural response patterns and low-level visual functionality (e.g. spatial resolution, contrast sensitivity, and orientation selectivity) are shaped by the statistics of the postnatal visual experience. Neural tuning in the visual cortex changes over the course of typical development (18–23), and moreover, atypical visual experience can permanently alter neural responses (24–36). Additional evidence comes from clinical observations of human infants with visual disorders such as cataracts, amblyopia, and strabismus. When these disorders are not treated, the resulting disruption in early visual experience can lead to permanent functional deficits (37–40).

Although studies have not yet revealed whether a cardinal orientation bias is present in typical early visual experience, the literature on the development of the oblique effect has found that infant perceptual performance is biased toward the cardinal orientations. The oblique effect can be reliably detected at 3 to 4 mo of age (41), although a weak oblique effect may be detectable at 6 wk of age (42). The strength of the effect then increases

Significance

Visual experience during infancy is crucial for the development of the human visual system. By analyzing head-mounted camera videos, we computed the distribution of orientations in infants' view of the world. We found that infants, even at just three weeks of age, experience a higher prevalence of horizontal and vertical orientations in their visual environments. Thus, an infant's visual diet is biased even before they develop the motor control necessary to consistently align themselves with gravity. These findings reveal the structure of low-level visual information in the infant visual diet, a crucial step in understanding how visual experience shapes sensory, perceptual, and cognitive development.

Author affiliations: [a]School of Optometry, Indiana University, Bloomington, IN 47405; [b]Department of Computer Science, Indiana University, Bloomington, IN 47405; [c]Program in Cognitive Science, Indiana University, Bloomington, IN 47405; [d]Department of Psychological and Brain Sciences, Bloomington, IN 47405; and [e]Pudiyador Association for Community Empowerment, Chennai 600020, India

Author contributions: S.J., L.B.S., T.R.C., and K.B. designed research; Z.J.P., S.J., and K.B. performed research; Z.J.P. and K.B. analyzed data; and Z.J.P., L.B.S., T.R.C., and K.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: kbonnen@iu.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2421277122/-/DCSupplemental.

Published April 14, 2025.

over the first year of life and well into early childhood (41, 42). These findings are consistent with the hypothesis that infants' see horizontal and vertical orientations more frequently than oblique orientations—a bias in their visual experience which may shape the development of their neural circuitry and thereby their visual function.

To assess infants' visual diets and test for the presence of a cardinal orientation bias, we gathered first-person egocentric video recordings showing the visual scene in front of infants ($n = 75$, aged 3 wk to 12 mo). The infants wore these cameras in their home environment while performing daily activities (Fig. 1*A*), helping to ensure that we sampled their typical visual experience. As our goal was to measure general patterns in visual input over the course of human development, our sample included infants of various age groups and from two very different communities: a Midwestern college town in the United States and an urban fishing village in southeastern India. Analysis of these recordings reveals that horizontal and vertical structure are indeed prevalent in early visual experience. This finding is crucial because it establishes that it is possible that the visual experience of the infant could be driving the development of the oblique effect in early infancy.

Results

We characterized infants' visual input by computing orientation distributions in images collected with head-mounted video cameras over long periods of natural behavior in infants' home environments ($n = 75$, aged 3 wk to 12 mo; Fig. 1). To ensure that our results were not specific to a particular geographic or cultural context, we included data from two notably different communities (Table 1): one in Indiana, USA ($n = 50$, 238 h of video), and one in Tamil Nadu, India ($n = 25$, 90 h of video). Together, the data consisted of 328 h of video (35 million video frames).

Because the visual experience of an infant changes rapidly during the first year of life, we sought to analyze the effects over time. To achieve this in our analyses, we divided the data into five groups based on infants' age at the time of recording: 3 to 12 wk, 13 to 24 wk, 25 to 36 wk, 37 to 48 wk, and 49 to 52 wk. Additionally, we used a temporal sampling approach, grouping video frames in 15-min samples, and averaging across the frames in each sample.

For each video frame, we measured the orientation of edges using the method published in ref. 7 (Fig. 1 *B–D*; see *Materials and Methods* for more details). That computation operates on multiple spatial scales (i.e. spatial frequency bands), producing an orientation histogram per spatial scale. We found similar results at each of our six spatial scales (centered at 0.06 to 2 cycles/degree; *SI Appendix*, Fig. S1). We also repeated the analysis on a per-frame basis (*SI Appendix*, Fig. S2). The range of spatial frequency information is limited by the camera specifications (i.e. field of view, pixel resolution; see *Materials and methods*). The spatial frequency bands we were able to measure are well matched to the spatial frequencies visible to infants. Compared to adult vision, infant vision is more limited in the range of visible spatial frequencies. Evidence from preferential looking studies suggests that visual acuity in 1-mo-old infants is approximately 2 cycles/degree (43), while evidence from VEP studies suggests that 5-wk-old infants have a visual acuity of around 5 cycles/degree (44). Studies of their contrast sensitivity functions suggest that there is peak sensitivity in lower spatial frequencies (0.25 to 1 cycle/degree)(44). For this reason, the figures in the main text show data from the spatial scale centered at 0.5 cycles/degree.

Predominance of Cardinal Orientations. Orientation histograms reveal very clear peaks at the two cardinal orientations for all but the very youngest group of infants (Fig. 2*A*). To quantify this orientation imbalance, we calculated the ratio of orientations near the cardinal axes (0°/90°) to those near the ordinal (i.e., intercardinal; 45°/135°) axes for all 15-min video samples (Fig. 2*B*). Ratios greater than 1 indicate an overrepresentation of cardinal orientations relative to ordinal, and 90% of video samples had a ratio greater than 1. The median ratio for all samples was 1.90 with a bootstrapped 95% CI of 1.81 to 2.03.

Looking across age groups, we found that the overrepresentation of cardinal orientations emerges early in infancy and that the strength of the overrepresentation changes over development. Across age, the ratios were well fit by a logarithmic curve (Fig. 2*B*, red line; $P < 0.01$), indicating that the ratio increases over development, with the most rapid increase occurring early in infancy. Indeed, the percentage of segments with a ratio greater than 1 increases over the first 24 wk, stabilizing very close to 100% after 25 wk (Fig. 2*C*). We also computed the median ratio per age group and, unsurprisingly, found that the median for the first age group (3 to 12 wk) is quite low compared to the older groups (Fig. 2*D*). Even so, the median ratios for all age groups were above 1, confirming that the cardinal orientation bias was a feature of visual experience for all ages in our sample.

Comparison of Infants in India and the United States. Our data collection included infants from two distinct communities: a mid-size college town in Indiana, USA, and an urban fishing village in Tamil Nadu, India. Approximately 73% of recordings came from the United States and 23% from India. In the United States, recordings were made primarily inside of infants'

Fig. 1. Measuring the distribution of orientations in egocentric video data. (*A*) Recording egocentric video in home environments. Two infant participants are shown wearing a lightweight Looxcie2 camera, secured to the head using a hat. (*B*) Example video frame from an outdoor environment. Our data consists of egocentric video frames, collected during infants' daily life. (*C*) Extracting orientation from video frames. To extract orientations from the example image in (*B*), we select a circular patch at the center of the video frame, identify locations with pronounced edges, and compute the orientation of those edges. (*D*) Measuring orientation frequency. We calculate orientation frequency histograms per-image by measuring the frequency of edges with particular orientations.

Table 1. Population and home environment

	Chennai, TN, India	Bloomington, IN, USA
Population density (persons/sq. km.)	26,553	1,340
Average people per home	4.5	2.2
Average dwelling size (sq. ft.)	223	1,742
Dwelling density (persons/1,000 sq. ft.)	20.1	1.3
Rooms per home	1.5	5.8

homes, and in India, recordings were made primarily outside of the home, within the home's urban surroundings. The recordings from both sites were dominated by man-made elements.

We compared the orientation distributions collected from American infants (n = 50) and Indian infants (n = 25) and found no significant differences (Fig. 3). To examine the effects of age and geographic location on the ratios, we performed a two-way ANOVA. Results showed no significant interaction between age and location ($P = 0.70$), and no main effect of location ($P = 0.15$) or age ($P = 0.09$).

In Fig. 3*C*, there is one outlying curve with peaks near the intercardinal orientations. This curve shows data for the youngest age group from Tamil Nadu, which had only five video segments. Four of the five segments were recorded by one 11-week-old infant. Visual inspection of that infant's data reveals that a large section of the recording consisted primarily of a ceiling view with well-defined ∼45° lines. That uniform segment of data resulted in the atypical pattern that can be seen in the outlying curve.

Discussion

In this study, we sought to determine whether the overrepresentation of horizontal and vertical structures in the world results in a biased distribution of orientation structure in the early visual experience. To measure the structure of infants' visual diet, we analyzed egocentric videos recorded from 75 infants ranging in age from 3 to 52 wk, living in either the United States or India.

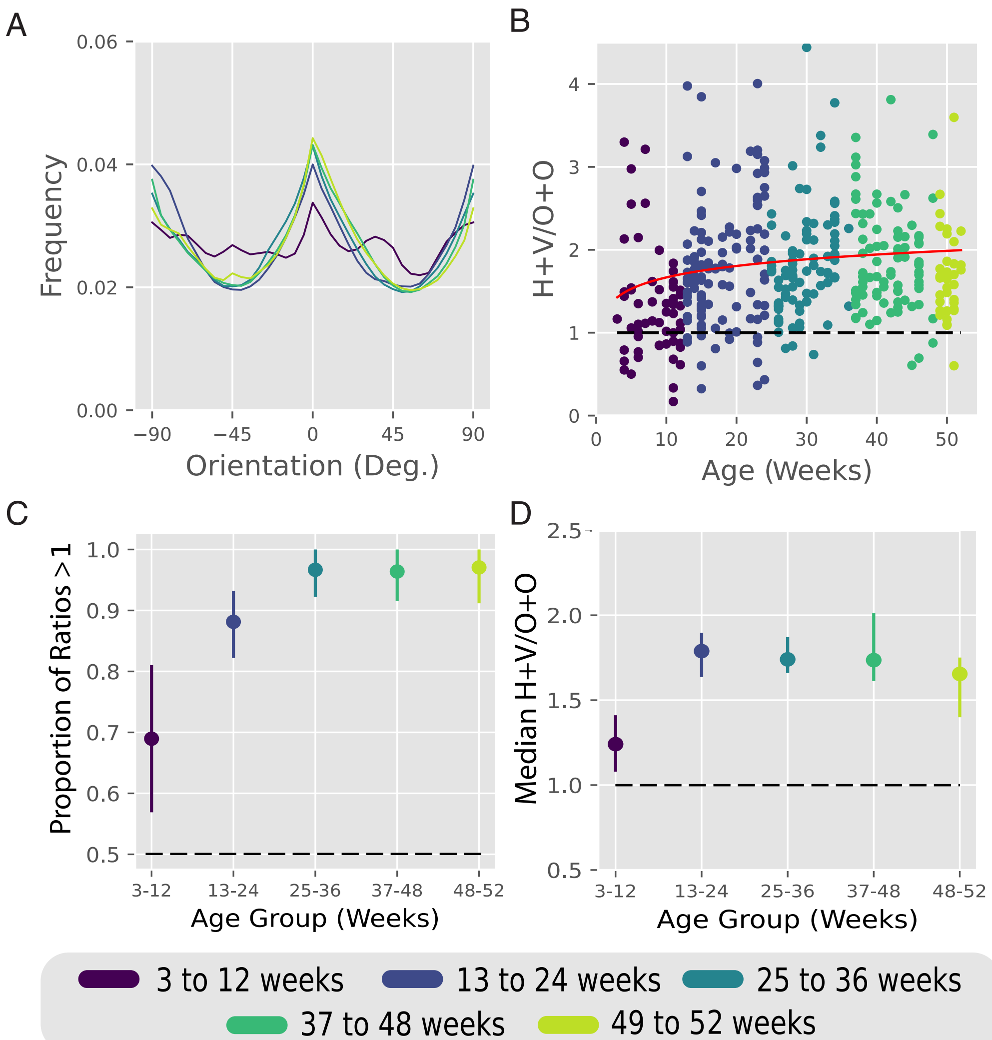

Fig. 2. Histograms of orientation show the predominance of cardinal orientations throughout early development. Results are shown for the spatial frequency band centered at 0.5 cycles/degree. Results for the other bands we analyzed are similar (*SI Appendix*, Fig. S1). (*A*) Average orientation histograms. Peaks occur at the cardinal orientations, even in the youngest age group. (Data are binned by age, with each color indicating 12-wk age bin.) (*B*) Ratio in prevalence of cardinal (H+V) to oblique orientations ($O_{45} + O_{135}$). The ratio of each video segment is plotted as a function of the infant's age, and the logarithmic line of best fit is plotted in red. Nearly all ratios (90%) are greater than 1, indicating cardinal orientations are more prevalent. (*C*) Proportion of ratios greater than 1. Plotted for each age group, with vertical lines indicating bootstrapped 95% CIs. The percentage greater than one is well above chance (50%) for all age groups, with a rising trend in the dominance of cardinal orientations as infants develop from 3 wk to 6 mo of age. (*D*) Median ratios. Plotted for each age group, with vertical lines indicating bootstrapped 95% CIs. The median ratios were all above 1 and follow a similar pattern to (*C*).

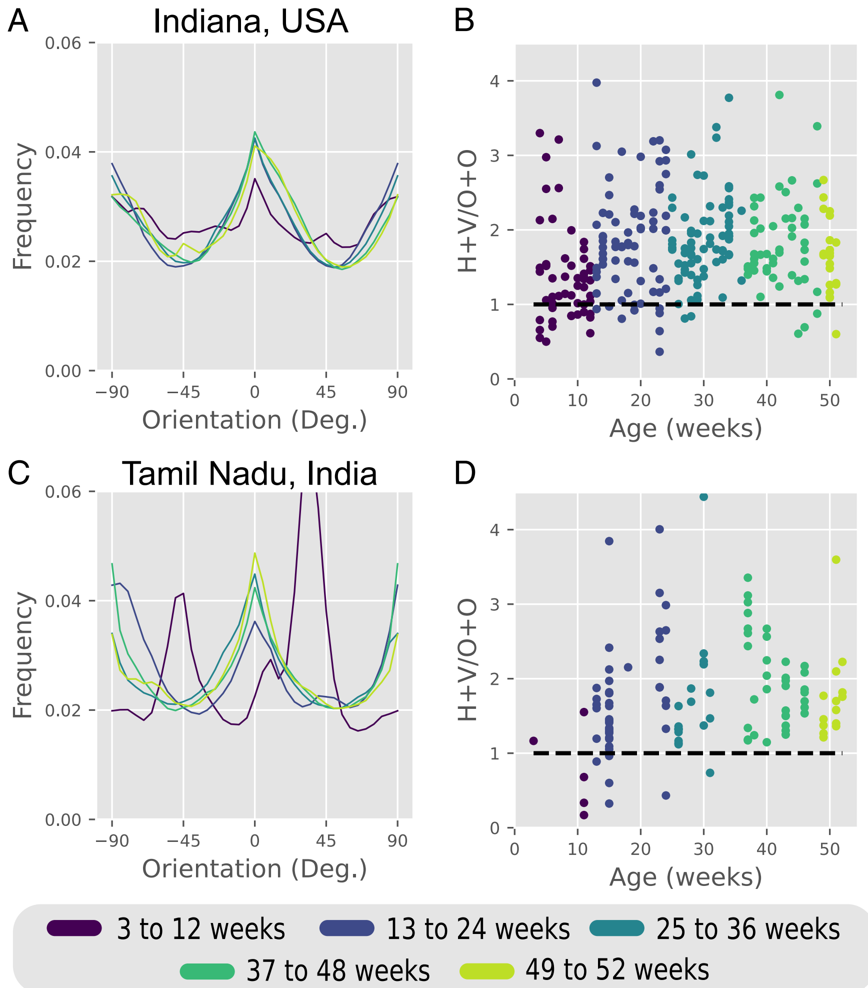

Fig. 3. Orientation histograms and cardinal/oblique ratios by site. (*A* and *B*) show data from infants living in Indiana, USA, while (*C* and *D*) show data from infants living in Tamil Nadu, India. (*A* or *C*) Average orientation histograms. Data are binned by age, with color indicating age by 12-wk bin. (*B* or *D*) Ratio in prevalence of cardinal (H+V) to ordinal orientations ($O_{45} + O_{135}$). Plotted per video segment as a function of infant age.

One limitation of this work was the different sample sizes between the two communities where the data was collected. Of the 328 h of video frames analyzed, 238 h came from Indiana (73%, 50 infants), and 90 h (27%, 25 infants) came from Tamil Nadu.

We found an overrepresentation of cardinal orientations in the captured visual scenes at all ages and in both locations. As the cardinal bias in the world's structure is relative to gravity, this finding suggests that even the youngest infants' head positions tend to be aligned with gravity. Furthermore, we found that the strength of the cardinal orientation bias increased from 3 to 12 wk, consistent with the development of a basic ability to orient the head to gravity during the first 12 wk of life (45). Because this bias was present in the data from both sites despite notable differences between recorded scenes, we expect the distribution of edge orientations to be biased from birth in infants from most countries and cultures.

The Changing Infant Experience. World scene statistics exhibit a cardinal bias relative to gravity, such that the orientations parallel and perpendicular to the direction of gravity (vertical and horizontal, respectively) are the most common (5, 7, 16). While this cardinal bias in the world results in a cardinal bias in adult visual input, infants have limited mobility and thus a limited ability to align their body with gravity. Because of this difficulty aligning to gravity, the cardinal bias could be absent from early infant visual experience. Such a lack of visual input bias would imply that the oblique effect observed in infant perceptual experiments must be more innate than learned. However, we found that infant visual experience does show a cardinal bias, even in young infants with limited motor control. Thus, infants are sufficiently aligned with gravity to experience the cardinal bias. Due to their lack of postural control, that alignment is likely due to caregiver placement. This notion is supported by evidence that infants' posture and resulting visual experience depend both on their developing motor skills and on caregiver placement (46–48). Furthermore, there is evidence that caregivers align infants with gravity, as caregiver hand placement on an infant's trunk is correlated with the infant's current trunk control capabilities (49).

Interestingly, we also found that the strength of the cardinal bias increases over the same period that motor control for head stability rapidly improves (45). This hints at a possible tie between visual and motor development. Studies show that early visual input calibrates proprioceptive and vestibular systems (50) and even that infants learning to stand use visual proprioceptive information more than nonvisual information (51). As such, the statistics of the infant visual experience may be influencing both the development of neural circuitry for low-level visual processing and providing cues that infants learn to use for postural stability.

While our findings suggest that the infant visual experience is somewhat similar to the adult visual experience in terms of the presence of a cardinal bias, we would like to emphasize that infants' view of the world is markedly different than that of adults in other ways. Studies have shown that infants' experiences of objects, faces, and hands are quite different from their caretakers' experiences and that they change markedly throughout the first 2 y of life (52, 53). For example, the relative frequency of faces compared to hands decreases over the first two years of life, with the youngest infants experiencing faces at much higher rates than the oldest infants (53). Additionally, there is evidence that the perceptual bias for cardinal orientations is retinally aligned in infants up to 7 mo of age (54) but gravitationally aligned in adults (55). In other words, an infant's perceptual performance depends on the orientation projected onto their retina, whereas an adult's perceptual performance depends more on the orientation in the world's reference frame, regardless of the orientation of their eyes or head.

Consistency of Statistics Across Scenes. In addition to finding that the cardinal bias is present for infants across the first year of life, we also found that the cardinal bias is present for infants in two distinct communities. This result is perhaps unsurprising given that the existence of a cardinal bias in the world's structure is very well-established. Previous studies have noted an overrepresentation of cardinal orientations in indoor, outdoor, and fully natural scenes (i.e., scenes containing no man-made components) (5, 7, 16). Though one may expect scenes with man-made components to have more horizontal and vertical structure than natural scenes, existing work suggests that the cardinal bias is present at similar or even greater levels in typical "pastoral" natural scenes, compared to both indoor and outdoor "carpentered" scenes with man-made components (5). Thus it is reasonable to expect that the cardinal orientation bias is present for most infants (and adults) in most communities.

However, there is perceptual evidence that the cardinal bias is not universal. Studies have shown that individuals from some cultures do not exhibit a strong oblique effect (56, 57). Typically, the "oblique effect" results in better performance on perceptual tasks (e.g., detection, discrimination) at horizontal and vertical orientations than at oblique orientations (6). However, these performance asymmetries are thought to be related to the existence of a cardinal bias in visual input, and individuals in some cultures do not experience a notable cardinal bias in their visual experience. As an example, individuals from an indigenous Cree community in Quebec, Canada, lacked a statistically significant oblique effect (56). Interestingly, the man-made structures in this community have contours at varied orientations, potentially resulting in a more uniform orientation distribution than would be found in carpentered scenes in other communities. Thus, one possibility is that the man-made elements of their environment have resulted in less biased visual input statistics and thereby less biased performance. An interesting avenue for future work could be the exploration of when and how the cardinal bias in visual input is eliminated in such communities.

The Egocentric Viewpoint. Studies of the cardinal bias in natural scene statistics have typically analyzed still photographs taken with the intention of representing a fixed adult viewpoint (5, 7, 16, 58), whereas we analyzed frames from egocentric video streams representing an infant viewpoint. Since egocentric videos capture an individual's view over time, the frames better sample an individual's everyday visual experience. Therefore, we are able to obtain more accurate estimates of the statistics of visual input.

Though, to be precise, we measured the orientation distribution for the head-centered view, as this dataset did not include gaze measurements. Obtaining a more direct measurement of the visual input would require that one account for eye movements. However, studies have shown that the head-centered view should be a relatively good approximation of the eye-centered view across the lifespan because 80% of the distribution of gaze in freely moving infants, toddlers, children, and adults falls within 20°of the center of the head's field of view (59). Thus, with sufficient sample size, the statistics of head camera images can reasonably estimate the statistics of visual input. This is particularly true for very young infants as we know that young infants rarely direct gaze to the periphery (60) and that infants as young as 2 wk postbirth turn their heads to look with eyes and head pointed in the same direction (60–62).

Learning Biases in the Visual Input. While existing evidence suggests that the neural and perceptual bias in favor of cardinal orientations is more learned than innate, a crucial missing piece of that puzzle has been whether or not early visual experience exhibits a cardinal orientation bias. Our findings provide evidence that infants do see cardinal orientations most frequently. Thus, the neural and perceptual cardinal bias may indeed be more learned than innate.

Results about the establishment of cardinal biases in neural activity and perceptual responses have a variety of potential applications. In particular, such results could be used to improve the training of artificial agents. Implementing a systematic change in visual experience over training, like the one we observe here, may allow for the development of artificial intelligence systems that require less training or learn to process visual input more efficiently.

Materials and Methods

Egocentric Video Data. In total, 328 h of video data (i.e., 35 million video frames) were collected by 75 infants, living in two distinct communities. We used a lightweight, battery-powered Looxcie2 camera that was secured to the infant's head using a snug fitting hat (Fig. 1*A*). The camera recorded video at 30 frames per second, with a pixel resolution of 480p (480 $\times$ 640 pixels) and a diagonal field of view of 62° (39° $\times$ 51°; analysis is performed on the central 256 $\times$ 256 pixels, the central 20 degrees). Parents were instructed on how to use the camera, given two hat-camera systems to take home, and asked to record 4 to 6 h of video during the infants' waking hours. We used the same equipment and gave parents the same instructions in both geographical locations.

These data were collected according to principles of the Declaration of Helsinki under IRB-approved protocols. The procedures for recruiting, consenting, data collection, and data management for both the US and Indian populations were reviewed and approved by the Institutional Review Board at Indiana University, and conducted under Protocol no. 1505862312 for the US data collection and Protocol no. 0909000021 for the India data collection. The U.S. families were recruited from an opt-in data base of families interested in participating in research and from outreach events in the community for families of young children. Parents were contacted through mail or email about the methods and goal of this study. If they were interested in participating, an appointment was arranged for training in the use of the head camera and consent was obtained. In India, the families were recruited with the assistance of the Pudiyador Community Organization, an education-driven private and family organization with programs for parents, children, and teens. Families were informed about the methods and goals of the project and if interested were visited by a member of the staff and experimenter (an Indian native and speaker of the local language) and were trained in the use of the head camera, and consent was obtained at that meeting.

Indiana, USA. In the United States, we collected approximately 238 h of egocentric video from 50 infants during their first year of life. The infants belonged to predominantly working- and middle-class families living in or near Monroe County. Monroe County contains the city of Bloomington, which is a mid-size college town. The families typically lived in homes with multiple large rooms, and they spent much of their daily life inside the house (i.e., within man-made environments). These data are included in the analysis published in ref. 53.

Tamil Nadu, India. In India, we collected approximately 90 h of egocentric video from 25 infants during their first year of life. The infants belonged to predominantly working-class families living in an urban fishing village in Tamil Nadu, near the city of Chennai in southeastern India. Families were recruited by the Pudiyador Association for Community Empowerment. They typically lived in densely clustered single-room homes. Most of their activities (including cooking, eating, and cleaning) occurred in outdoor urban environments (i.e., scenes included many man-made elements).

Initial Video Processing. From each 640×480 video frame, we extracted the central 256×256 pixel region. We then converted to gray-scale using the ITU-R 601-2 luma transform via the OpenCV Python library (63) and normalized the pixel intensities to values between zero and one.

Filtering frames based on motion. Blur from motion affects the low-level statistics of an image, introducing false apparent edges. Because rapid head movements can create these motion artifacts, we excluded video frames recorded during head movements from our analysis. To detect infants' head movements, we computed the camera movement between each sequential pair of frames (64). If the camera moved 10°/s or faster, we excluded the frame.

Measuring Orientation Histograms. To measure orientation histograms for each image, we used the method outlined in ref. 7. This analysis, which we will refer to as the Girshick method, extracts the edges in a given image and measures the orientations of those edges. In this section, we describe the implementation of this method. Briefly, the Girshick method creates multiple versions of an image, each with a different spatial scale; quantifies the local structure at each point in each image version; identifies all points with sufficient contrast and orientation strength; computes the dominant orientation at those locations; and outputs a list of edges and their orientations per spatial scale. That information is then used to compute orientation histograms, which represent the relative frequency of each orientation in an image at a particular spatial scale.

Multiscale image analysis. The Girshick method creates a Gaussian pyramid, which is a multiscale representation of the image. The pyramid's layers are produced by repeatedly applying Gaussian blurring, followed by down-sampling. As a result, each layer contains progressively lower spatial frequency details, allowing for the analysis of the image at different levels of spatial detail. The remaining analyses are performed on each of the spatial scales in the pyramid.

Quantifying local structure. To compute the orientation structure in the neighborhood surrounding each point in an image, the Girshick method computes a structure tensor per point. A structure tensor (also called a second-moment-matrix) is a 2×2 matrix that summarizes the orientation structure of a neighborhood around a specified point.

For each pixel location (i, j) in an image I, the structure tensor, S, is given by

$$S(i,j) = \begin{bmatrix} s_{xx}(i,j) & s_{xy}(i,j) \\ s_{xy}(i,j) & s_{yy}(i,j) \end{bmatrix},$$

where s_{xx}, s_{yy}, and s_{xy} are smoothed products of the image gradients ∇I_x and ∇I_y.

$$s_{xx}(i,j) = \sum_{m=-1}^{1} \sum_{n=-1}^{1} \nabla I_x(i+m, j+n)^2 \cdot B(m+1, n+1)$$

$$s_{yy}(i,j) = \sum_{m=-1}^{1} \sum_{n=-1}^{1} \nabla I_y(i+m, j+n)^2 \cdot B(m+1, n+1)$$

$$s_{xy}(i,j) = \sum_{m=-1}^{1} \sum_{n=-1}^{1} \nabla I_x(i+m, j+n) \cdot \nabla I_y(i+m, j+n) \cdot B(m+1, n+1),$$

where ∇I_x is the image gradient in the x-direction, ∇I_y is the image gradient in the y-direction, and B is a smoothing kernel. The matrix B is defined as

$$B = \frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

The two image gradients, ∇I_x and ∇I_y, quantify the rate of intensity change at each pixel location (i, j) in the image I:

$$\nabla I_x(i,j) = \sum_{m=-2}^{2} \sum_{n=-2}^{2} I(i+m, j+n) \cdot G_x(m+2, n+2)$$

$$\nabla I_y(i,j) = \sum_{m=-2}^{2} \sum_{n=-2}^{2} I(i+m, j+n) \cdot G_y(m+2, n+2),$$

where G_x and G_y are direction-specific derivative filters. The filters G are the tensor product of two component 5-point kernels: derivative kernel, g_d, and smoothing kernel, g_s:

$$G_x = g_d \otimes g_s$$
$$G_y = g_s \otimes g_d.$$

For the Girshick method, these kernels are defined as

$$g_s = [0.0375659, 0.249153, 0.426375, 0.249153, 0.0375659]$$
$$g_d = [-0.109604, -0.276691, 0.0, 0.276691, 0.109604].$$

Notably, the application of these derivative kernels is what limits the measurements of orientation to specific spatial frequency bands. The blurring/down-sampling process creates "spatial scale images," progressively removing high frequency content. The size of the "spatial scale image" determines the upper bound of the spatial frequency band and the size of the 5×5 derivative kernel (5 pixels) determines the lower bound of the spatial frequency band. The highest and lowest spatial frequency information available across all spatial frequency bands is determined by the field of view and pixel resolution of the camera.

These computations produce a structure tensor S for each layer of the multiscale representation (i.e., Gaussian pyramid) of each video frame. Each layer's tensor is composed of 2×2 matrices $S(i, j)$ that summarize the orientation structure for the neighborhood around every point in that layer.

Finding edges. A structure tensor can be characterized by computing its eigenvalues. Generally, an eigenvalue λ (also known as a characteristic value or characteristic root) of a linear transformation matrix T represents the scaling of eigenvector $\mathbf{v}$ that occurs when the linear transformation is applied to that vector: $T\mathbf{v} = \lambda\mathbf{v}$. The exact significance of a matrix's eigenvalues depends on the context. In the context of a structure tensor, the eigenvalues represent the strength of the image gradients at the image location in question. If one eigenvalue is significantly larger than the other, that would indicate a strong intensity gradient, signifying the presence of a high-contrast oriented line.

The two eigenvalues, λ_1 and λ_2, of structure tensor S are

$$\lambda_1(i,j) = \alpha(i,j) + \beta(i,j)$$
$$\lambda_2(i,j) = \alpha(i,j) - \beta(i,j),$$

where α and β are intermediary terms that depend on matrices s_{xx}, s_{yy}, and s_{xy}. Specifically:

$$\alpha(i,j) = \frac{s_{xx}(i,j) + s_{yy}(i,j)}{2}$$
$$\beta(i,j) = \sqrt{\alpha(i,j)^2 - (s_{xx}(i,j) \cdot s_{yy}(i,j) - s_{xy}(i,j)^2)}.$$

With the eigenvalues of a tensor, we can calculate both contrast energy, c, and orientedness, r. Those values describe image properties near a particular point (i, j) that we use to determine whether an oriented edge is present at that location. The contrast energy, c, is the sum of the eigenvalues:

$$c(i,j) = \lambda_1(i,j) + \lambda_2(i,j)$$

The orientedness (i.e., the strength of the most prominent orientation), r, can be calculated as the relative difference squared between $\lambda_1(i,j)$ and $\lambda_2(i,j)$:

$$r(i,j) = \left(\frac{\lambda_1(i,j) - \lambda_2(i,j)}{\lambda_1(i,j) + \lambda_2(i,j) + \varepsilon}\right)^2,$$

where ε is a small constant that prevents division by zero ($\varepsilon = 2.2204$ e −16).

A point (i,j) is said to contain an oriented edge if both its contrast energy $c(i,j)$ and orientedness $r(i,j)$ are above the thresholds t_c and t_r, respectively. The orientedness threshold, t_r, is an absolute threshold of 0.8. Meanwhile, the contrast threshold, t_c, is a relative threshold set to the 68^{th} percentile of all the values c in a layer of the multiscale image representation, with a minimum threshold of 1.0 e −4.

Computing orientation. To determine the orientation, o, at each location, we use the equation:

$$o(i,j) = \arctan\left(\frac{s_{xx}(i,j) - \lambda_2(i,j)}{s_{xy}(i,j)}\right) - \frac{\pi}{2}.$$

Because there may not be a dominant orientation or sufficient contrast to perceive an edge at every location, the next step is to apply the contrast and orientedness thresholds, t_c and t_r. Only orientations at locations where there is both sufficient contrast ($c(i,j) > t_c$) and a sufficiently dominant orientation ($r(i,j) > t_r$) are included in further analyses; other values are removed from o.

Finally, to remove any bias that could be introduced by artifacts produced at image boundaries, the method applies a circular mask to the orientation matrix o. The mask's radius was 90% of the cropped image's height and width. Orientation values, $o(i,j)$, were excluded from further analyses if they fell outside that circular mask.

Computing the orientation histogram. The orientation histogram is calculated by binning by every 5° from −87.5° to 87.5° and then creating an additional 5° bin for orientation values between ±87.5°. We then normalize the histogram by dividing each bin's count by the total number of pixel locations belonging to oriented edges. The normalized histogram provides a probabilistic representation, indicating the likelihood that an edge within the image exhibits an orientation from a specific bin.

Quantifying Orientation Bias. To quantify the relative frequency of horizontal and vertical orientations in infants' egocentric views, we calculated the ratio of cardinal orientations to ordinal (i.e., intercardinal) orientations $(H + V)/(O_{45} + O_{135})$. We first aggregated the orientation frequency in the angular region ±7.5° around each of the two cardinal orientations (horizontal and vertical; 0°/90°) and the two ordinal orientations (45°/135°). After calculating the cardinal orientation frequency $(H + V)$ and the ordinal orientation frequency $(O_{45} + O_{135})$, we then divided to compute the ratio of cardinal orientations to ordinal orientations. A ratio of 1 would indicate equal cardinal and ordinal orientation content, while a ratio greater than 1 would indicate that cardinal orientations were more prevalent than ordinal orientations (and vice versa for a ratio less than 1).

Data, Materials, and Software Availability. Some study data are available: Data analyses were performed on a video dataset that cannot be shared publicly for privacy reasons under the terms of the IRB and the NSF/NIH grants that supported the data collection. We consider the per-image orientation histograms as the raw measurements for this paper. We will share the orientation data and the software to perform the analyses presented in this paper.

ACKNOWLEDGMENTS. We would like to thank the families who collected the data for this study. Additionally, we thank the Pudiyador Association for Community Empowerment for their support in Chennai, India. We thank Ahna Girshick for sharing an implementation of the algorithm originally used in their study. We thank Stephanie M. Shields for thoughtful comments and assistance in revising the manuscript. The research reported in this publication was supported by the NIH NEI, under award number R01EY032897 to L.S. and T.R.C.; the NSF, under award number BCS152398 to L.S.; and also in part by the Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute.

1. T. A. Knight, V. On the direction of the radicle and germen during the vegetation of seeds. By Thomas Andrew Knight, Esq. F. R. S. In a letter to the Right Hon. Sir Joseph Banks, K. B. P. R. S. *Philos. Trans. R. Soc. Lond.* **96**, 99–108 (1806).
2. C. J. Dakin, A. Rosenberg, "Chapter 3 - Gravity estimation and verticality perception" in *Handbook of Clinical Neurology, Balance, Gait, and Falls*, B. L. Day, S. R. Lord, Eds. (Elsevier, 2018), vol. 159, pp. 43–59.
3. D. P. Häder, M. Braun, R. Hemmersbach, "Gravity sensing, graviorientation and microgravity" in *Gravitational Biology I: Gravity Sensing and Graviorientation in Microorganisms and Plants*, M. Braun, M. Böhmer, D. P. Häder, R. Hemmersbach, K. Palme, Eds. (Springer International Publishing, Cham, 2018), pp. 1–11.
4. R. H. Anken, H. Rahmann, "Gravitational zoology: How animals use and cope with gravity" in *Astrobiology: The Quest for the Conditions of Life*, G. Horneck, C. Baumstark-Khan, Eds. (Springer, Berlin, Heidelberg, 2002), pp. 315–333.
5. E. Switkes, M. J. Mayer, J. A. Sloan, Spatial frequency analysis of the visual environment: Anisotropy and the carpentered environment hypothesis. *Vis. Res.* **18**, 1393–1399 (1978).
6. S. Appelle, Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals: Psychological Bulletin. *Psychol. Bull.* **78**, 266–278 (1972).
7. A. R. Girshick, M. S. Landy, E. P. Simoncelli, Cardinal rules: Visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* **14**, 926–932 (2011).
8. E. Gentaz, G. Baud-Bovy, M. Luyat, The haptic perception of spatial orientations. *Exp. Brain Res.* **187**, 331–348 (2008).
9. P. Mustillo, E. Francis, S. Oross, R. Fox, G. A. Orban, Anisotropies in global stereoscopic orientation discrimination. *Vis. Res.* **28**, 1315–1321 (1988).
10. N. Smyrnis, A. Mantas, I. Evdokimidis, "Motor oblique effect": Perceptual direction discrimination and pointing to memorized visual targets share the same preference for cardinal orientations. *J. Neurophysiol.* **97**, 1068–1077 (2007).
11. R. J. W. Mansfield, Neural basis of orientation perception in primate vision. *Science* **186**, 1133–1135 (1974).
12. C. S. Furmanski, S. A. Engel, An oblique effect in human primary visual cortex. *Nat. Neurosci.* **3**, 535–536 (2000).
13. B. Li, M. R. Peterson, R. D. Freeman, Oblique effect: A neural basis in the visual cortex. *J. Neurophysiol.* **90**, 204–217 (2003).
14. R. L. Boltz, R. S. Harwerth, E. L. Smith, Orientation anisotropy of visual stimuli in rhesus monkey: A behavioral study. *Science* **205**, 511–513 (1979).
15. A. van der Schaaf, J. H. van Hateren, Modelling the power spectra of natural images: Statistics and information. *Vis. Res.* **36**, 2759–2770 (1996).
16. D. M. Coppola, H. R. Purves, A. N. McCoy, D. Purves, The distribution of oriented contours in the real world. *Proc. Natl. Acad. Sci. U.S.A.* **95**, 4002–4006 (1998).
17. M. Henderson, J. T. Serences, Biased orientation representations can be explained by experience with nonuniform training set statistics. *J. Vis.* **21**, 10 (2021).
18. D. H. Hubel, T. N. Wiesel, Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *J. Neurophysiol.* **26**, 994–1002 (1963).
19. B. Chapman, T. Bonhoeffer, Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. *Proc. Natl. Acad. Sci. U.S.A.* **95**, 2609–2614 (1998).
20. Y. Li, D. Fitzpatrick, L. E. White, The development of direction selectivity in ferret visual cortex requires early visual experience. *Nat. Neurosci.* **9**, 676–681 (2006).
21. G. Shen, X. Tao, B. Zhang, E. L. Smith III, Y. M. Chino, Oblique effect in visual area 2 of macaque monkeys. *J. Vis.* **14**, 3 (2014).
22. Y. Gu, J. Cang, Binocular matching of thalamocortical and intracortical circuits in the mouse visual cortex. *eLife* **5**, e22032 (2016).
23. J. T. Chang, D. Whitney, D. Fitzpatrick, Experience-dependent reorganization drives development of a binocularly unified cortical representation of orientation. *Neuron* **107**, 338–350.e5 (2020).
24. T. N. Wiesel, D. H. Hubel, Extent of recovery from the effects of visual deprivation in kittens. *J. Neurophysiol.* **28**, 1060–1072 (1965).
25. C. Blakemore, G. F. Cooper, Development of the brain depends on the visual environment. *Nature* **228**, 477–478 (1970).
26. A. Hein, R. Held, E. C. Gower, Development and segmentation of visually controlled movement by selective exposure during rearing. *J. Comp. Physiol. Psychol.* **73**, 181–187 (1970).
27. H. V. Hirsch, D. N. Spinelli, Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science* **168**, 869–871 (1970).
28. A. G. Leventhal, H. V. Hirsch, Cortical effect of early exposure to diagonal lines. *Science* **190**, 902–904 (1975).
29. M. P. Stryker, H. Sherk, Modification of cortical orientation selectivity in the cat by restricted visual experience: A reexamination. *Science* **190**, 904–906 (1975).
30. A. G. Leventhal, H. V. Hirsch, Effects of early experience upon orientation sensitivity and binocularity of neurons in visual cortex of cats. *Proc. Natl. Acad. Sci. U.S.A.* **74**, 1272–1276 (1977).
31. P. G. Shinkman, C. J. Bruce, Binocular differences in cortical receptive fields of kittens after rotationally disparate binocular experience. *Science* **197**, 285–287 (1977).
32. C. R. Olson, R. D. Freeman, Monocular deprivation and recovery during sensitive period in kittens. *J. Neurophysiol.* **41**, 65–74 (1978).
33. D. E. Mitchell, "The influence of early visual experience on visual perception" in *Visual Coding and Adaptability* (Psychology Press, 1981).
34. R. von der Heydt, P. Hänny, M. Dürsteier, "The role of orientation disparity in stereoscopic perception and the development of binocular correspondence" in *Sensory Functions, Advances in Physiological Sciences*, E. Grastyán, P. Molnár, Eds. (Pergamon, 1981), vol. 16, pp. 461–470.

35. S. L. Cloherty *et al*., Sensory experience modifies feature map relationships in visual cortex. *eLife* **5**, e13911 (2016).
36. J. N. Levine, H. Chen, Y. Gu, J. Cang, Environmental enrichment rescues binocular matching of orientation preference in the mouse visual cortex. *J. Neurosci. Off. J. Soc. Neurosci.* **37**, 5822–5833 (2017).
37. M. Von Senden, *Space and Sight: The Perception of Space and Shape in the Congenitally Blind before and after Operation* (Free Press of Glencoe, Glencoe, IL, 1960).
38. R. Sireteanu, W. Singer, The "vertical effect" in human squint amblyopia. *Exp. Brain Res.* **40**, 354–357 (1980).
39. M. J. Reed, M. J. Steinbach, H. Ono, S. Kraft, B. Gallie, Alignment ability of strabismic and eye enucleated subjects on the horizontal and oblique meridians. *Vis. Res.* **35**, 2523–2528 (1995).
40. D. H. Hubel, T. N. Wiesel, Early exploration of the visual cortex. *Neuron* **20**, 401–412 (1998).
41. J. Gwiazda, S. Brill, I. Mohindra, R. Held, Infant visual acuity and its meridional variation. *Vis. Res.* **18**, 1557–1564 (1978).
42. S. C. Leehey, A. Moskowitz-Cook, S. Brill, R. Held, Orientational anisotropy in infant vision. *Science* **190**, 900–902 (1975).
43. D. Y. Teller, The development of visual acuity in human and monkey infants. *Trends Neurosci.* **4**, 21–24 (1981).
44. A. M. Norcia, C. W. Tyler, R. D. Hamer, Development of contrast sensitivity in the human infant. *Vis. Res.* **30**, 1475–1486 (1990).
45. B. Bertenthal, C. von Hofsten, Eye, head and trunk control: The foundation for manual development. *Neurosci. Biobehav. Rev.* **22**, 515–520 (1998).
46. K. E. Adolph, J. E. Hoch, Motor development: Embodied, embedded, enculturated, and enabling. *Annu. Rev. Psychol.* **70**, 141–164 (2019).
47. J. M. Franchak, The ecology of infants' perceptual-motor exploration. *Curr. Opin. Psychol.* **32**, 110–114 (2020).
48. J. M. Franchak, L. Smith, C. Yu, Developmental changes in how head orientation structures infants' visual attention. *Dev. Psychobiol.* **66**, e22538 (2024).
49. K. Duncan, A. Goodworth, C. S. N. Da Costa, M. Wininger, S. Saavedra, Parent handling of typical infants varies segmentally across development of postural control. *Exp. Brain Res.* **236**, 645–654 (2018).
50. K. Hötting *et al*., Postural control depends on early visual experience. *J. Vis.* **24**, 3 (2024).
51. D. N. Lee, E. Aronson, Visual proprioceptive control of standing in human infants. *Percept. Psychophys.* **15**, 529–532 (1974).
52. S. Bambach, D. J. Crandall, L. B. Smith, C. Yu, "Active viewing in toddlers facilitates visual object learning: An egocentric vision approach" in *Proceedings of the Annual Meeting of the Cognitive Science Society* (2016), vol. 38, pp. 1631–1636.
53. C. M. Fausey, S. Jayaraman, L. B. Smith, From faces to hands: Changing visual input in the first two years. *Cognition* **152**, 101–107 (2016).
54. F. Jouen, The influence of body position on perception of orientation in infants. *Behav. Brain Res.* **15**, 241–245 (1985).
55. K. Mikellidou, G. M. Cicchini, P. G. Thompson, D. C. Burr, The oblique effect is both allocentric and egocentric. *J. Vis.* **15**, 24 (2015).
56. R. C. Annis, B. Frost, Human visual ecology and orientation anisotropies in acuity. *Science* **182**, 729–731 (1973).
57. B. N. Timney, D. W. Muir, Orientation anisotropy: Incidence and magnitude in caucasian and Chinese subjects. *Science* **193**, 699–701 (1976).
58. B. C. Hansen, E. A. Essock, A horizontal bias in human visual processing of orientation and its correspondence to the structural components of natural scenes. *J. Vis.* **4**, 5 (2004).
59. L. Smith, C. Yu, H. Yoshida, C. M. Fausey, Contributions of head-mounted cameras to studying the visual environments of infants and young children. *J. Cogn. Dev. Off. J. Cogn. Dev. Soc.* **16**, 407–419 (2015).
60. T. R. Candy *et al*., Infants' use of eye movements to explore their natural environment. *J. Vis.* **24**, 974 (2024).
61. D. M. Regal, D. H. Ashmead, P. Salapatek, The coordination of eye and head movements during early infancy: A selective review. *Behav. Brain Res.* **10**, 125–132 (1983).
62. H. Bloch, I. Carchon, On the onset of eye-head coordination in infants. *Behav. Brain Res.* **49**, 85–90 (1992).
63. G. Bradski, The OpenCV library. *Dr. Dobb's J. Softw. Tools* **120**, 122–125 (2000).
64. G. Yang, D. Ramanan, "Learning to segment rigid motions from two frames." in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (IEEE, Nashville, TN, 2021).