

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/201289>

Please be advised that this information was generated on 2019-10-20 and may be subject to change.



Original Articles

Infants differentially update their internal models of a dynamic environment

E. Kayhan^{a,b,*}, S. Hunnius^b, J.X. O'Reilly^c, H. Bekkering^b^a Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany^b Donders Institute for Brain, Cognition and Behavior, Radboud University Nijmegen, the Netherlands^c University of Oxford, Oxford Centre for Functional MRI of the Brain, Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, Oxford, UK

ARTICLE INFO

Keywords:

Surprise
Prediction
Model update
Development
Learning

ABSTRACT

Unexpected events provide us with opportunities for learning about what to expect from the world around us. Using a saccadic-planning paradigm, we investigated whether and how infants and adults represent the statistics of a changing environment (i.e. build an internal model of the environment). Participants observed differently colored bees that appeared at an unexpected location every few trials. The color cues indicated whether the subsequent bees would appear at this new location (i.e. update trials) or at the same location as previously (i.e. no-update trials). Infants learned the predictive value of the color cues and updated their internal models when necessary. Unlike infants, adults had a tendency to update their models each time they observed a change in the structure. We argue that infants are open to learning from current evidence due to being less influenced by their prior knowledge. This is an advantageous learning strategy to form accurate representations in dynamic environments, which is fundamental for successful adaptation.

1. Introduction

Learning is the process that allows organisms to construct and modulate an internal model of the outside world for efficient functioning (Friston, 2010; Picard & Friston, 2014). Surprising events challenge one's current model of the world, as they often demand adjustments in predictions generated by these models (O'Reilly, 2013; O'Reilly et al., 2013). Here, we used a saccadic planning task to investigate whether and how infants and adults form internal models that represent the statistics of a dynamic environment and adjust their models when confronted with changes.

Maintaining an accurate representation of the world is a challenge for learners especially in environments that change over time. Surprising or unexpected events signaling change render old information irrelevant to the present state of the world, often making one's previous models of the environment obsolete (Courville, Daw, & Touretzky, 2006). This would then promote more learning, as internal models would be adjusted to represent the outside world accurately (Courville et al., 2006; Payzan-LeNestour & Bossaerts, 2011).

O'Reilly et al. (2013) investigated how humans update their internal models in a changing environment and whether a normative Bayesian learner model could explain this process. Using fMRI and a saccadic planning paradigm, they presented adults with targets on a circle that changed location in every few trials. The first trial in which the targets

appeared at a new location were called “update trials”. After “update trials”, the targets continued to appear around the same area for a couple of trials (i.e. expected trials). In some of the trials, the targets appeared at a new location on the circle for once (i.e. one-off trials); however, in the following trials the targets were presented around the same area as in the immediately preceding trials. Although the sudden appearances of the targets at new locations were unexpected in both “update” and “one-off” trials, participants were required to update their models in the update trials, but not in one-off trials, as the targets continued to appear at this new location only after the update trials. Data revealed activation in the parietal cortex when an immediate motor response was programmed as the location of the targets were unexpectedly changed in both trial types. However, the anterior cingulate cortex was specifically active only when participants updated their internal models to accommodate the information about the new location of the targets. In summary, these findings show that when the uncertainty about the parameters of the environment increases, adults adjust their internal models to represent the statistics of the outside world accurately.

Similarly, for efficient functioning, developing systems should not only detect changes but also use this information to adapt their internal models of the world. There is ample evidence in the developmental literature showing that even human fetuses and newborn infants detect changes and respond to deviations from regularities (Draganova et al.,

* Corresponding author at: Max Planck Institute for Human Cognitive and Brain Sciences, Stephanstraße 1A, 04103 Leipzig, Germany.

E-mail address: kayhan@cbs.mpg.de (E. Kayhan).

2005; Partanen et al., 2013; Sambeth et al., 2009). However, whether and how developing systems go beyond change detection and use surprising information to learn needs further investigation.

Stahl and Feigenson (2015) explored surprise-driven learning in human infants. Using a behavioral paradigm, they showed that 11-month-old infants explored objects more often and learned new functions of the objects more rapidly, when these objects were presented previously in situations that disconfirmed their expectations about core knowledge in physics (i.e. knowledge about physical rules that is already present prior to learning). Although this study demonstrates that unexpected events enhance further learning, it is not clear how infants form expectations based on surprising information. Because core knowledge is assumed to be already present very early on in life, the important question of whether and how infants use unexpected information to generate internal models to make predictions remains unanswered. Do they, for example, use each unexpected information to build an internal model? Alternatively, do they dissociate between unexpected information in the extent to which they are relevant for modulating predictions?

Previous research has shown that an infant brain is capable of forming internal models to make predictions. In a functional near infrared spectroscopy (fNIRS) study using an omission paradigm, Emberson, Richards, and Aslin (2015), have provided the first evidence that the predictive architecture of the brain is already present as early as 6 months. Emberson et al. (2015) demonstrated that after a learning period, when images were unexpectedly omitted, infants showed activation in the occipital cortex, as if an image was presented, suggesting that they generated predictions about the visual input. Importantly, this activation was not observed, if omission was expected to happen. Using EEG with a cross-modal cueing paradigm, Kouider et al. (2015) provided further evidence that infants formed predictions as a result of learning associations between auditory cues and visual categories, and their neural responses were differentially regulated based on the prior knowledge they acquired. Whereas early components were amplified for valid cues, late components, such as positive slow wave (PSW), were enhanced for invalid targets. These studies together suggest that the infant brain is already capable of forming predictions based on prior knowledge and sensory activity is modulated by the violations of these predictions. Although these studies provide initial evidence on the predictive nature of the infant brain, it remains unknown whether infants, like adults, dynamically update their internal models in response to changes in the environment.

We used a saccadic planning paradigm, similar to O'Reilly et al. (2013), to address three fundamental points. First, we investigated whether and how infants and adults *construct* an internal model of a changing environment that would then allow them to make predictions. Second, we examined whether infants and adults *dissociate* between unexpected information in the extent to which they are relevant for updating predictions, as it would be inefficient to adjust predictions with respect to irrelevant changes. Third, we asked whether infants and adults use *different strategies* to generate and modulate their internal models based on surprising information. Using the same procedure (i.e. same paradigm, no instructions), we ensured that infants and adults had similar prior information about the task specifics. We also tested another group of adults with instructions.

2. Material and methods

2.1. Participants

Thirty-nine 14-month-olds ($M = 14$ months 6 days, range: 13 months, 23 days–14 months, 12 days, 18 girls) were recruited for the study based on power analysis estimating a medium effect size, as there was no previous study using the same paradigm in infants. We considered the potential drop out in the power analysis given that this phenomenon is commonly observed in infancy research. Because

previous studies show that by the end of their first year of life, infants advance in their statistical learning skills, we investigated infants at 14 months of age (Saffran & Kirkham, 2018, for a review). Three infants did not complete the eye-tracking calibration procedure and seven infants were excluded during data processing because of failing to reach the inclusion criteria (see data processing section). The final sample thus included twenty-nine infants ($M = 14$ months, 6 days, range: 13 months, 27 days–14 months, 12 days, 12 girls). Participants were recruited from a database of volunteer families. Families received a baby book or 10 Euros for their participation. All data from adult participants were included in the analyses. Adult participants were recruited from a database of university students ($M = 23.41$ years, $SD = 3.34$ years, range 19–33 years, 36 female, 13 male). Adult participants received gift cards of five Euros for their participation. The local social science faculty's ethical committee approved the study (approval number: ECG2012-0910-058 DCC-NWO-EUea-Bekkering and ECG2012-1301-006 Stapel/Hunnus).

2.2. Materials and procedure

The same stimuli were used for testing infants and adults. We used a saccadic planning task in which we manipulated color and location of target stimuli. Prior to the start of trials, we showed participants an illustration of three differently colored bees on a garden background for 3000 ms, as if the bees were flying outdoors. Each trial started with an image of target position holders on a circle. After 500 ms, a fixation image (i.e. a sun image) appeared in the middle of the screen, which remained visible for 1000 ms (see Fig. 1). Following this, participants observed differently colored bees that appeared on one of eight designated spots on the circle and stayed visible for 1500 ms. One trial lasted for a total duration of 3000 ms, and the experiment consisted of 126 trials in total. The trials followed each other continuously as a movie sequence rather than being presented as separate images. The entire stimulus presentation took 6 min and 18 s. An example stimulus video can be found in the [Supplementary Material](#).

The experiment consisted of 98 “expected” trials and 28 “surprise” trials. There were two types of surprise trials that occurred with equal frequencies: surprise/update and surprise/no-update trials. We used color cues to indicate different trials. Bees could appear in three different colors (see Fig. 1). We alternated between two colors to display surprise/update and expected trials. The surprise/update trial bee always had a different color from the ongoing expected trial series and indicated the start of a new series of expected trials of that color. The color of the surprise/no-update trial bee remained the same throughout the experiment.

We manipulated the location of the bees in order to dissociate surprise that demanded an update and surprise that did not demand an update in predictions. In all surprise trials, the location (and color) of the bee differed from the previous trials. Participants were thus required to perform an eye movement to a new bee location, whereas in expected trials, the bee location (and color) remained unchanged. The two sorts of surprise trials differed in the trials that followed them. The surprise/update trials were the first of a series of trials of the same kind whereas the surprise/no-update trials were a one-off event with the subsequent trials' bee locations (and color) returning to what they had been before. Therefore, after a surprise/update trial, the subsequent bees continued to appear at the same location for on average four trials and participants thus had to update their predictions about the location of the bees. In the surprise/no-update trials, the location of the bee was again surprising, as compared to the location of the previous one; however, as this was a one-time event, participants were expected to *not* change their predictions about the location of the following bees. In expected trials, the bee appeared where it appeared in the previous trial. Fig. 1 illustrates the paradigm.

Together with the fixation image, a brief sound was played, and with each bee image, a jump sound was played that slightly varied in



Fig. 1. Screen snapshots for the different stimulus types within a continuous sequence of trials.

pitch in order to keep infants’ attention on the stimuli. The colors of the bees and associated sounds were counterbalanced across participants. The trial order, target positions and the distance of the surprising targets from the previous targets were pseudo-randomized.

The stimulus materials were created using Psychopy (version 1.83.04) which were converted into movies in 1920 × 1080 screen resolution using the open source software Kazaam (version 1.4.5). The stimuli were presented using Tobii Studio Software (3.3.0). Gaze data were recorded at 120 Hz by a corneal reflection eye tracker (Tobii TX300, Tobii Technology, Danderyd, Sweden) calibrated using a 9-point procedure. The procedure was repeated if only seven or fewer

calibration points were detected. Adults were seated on a chair in front of the eye-tracker screen and infants were seated in a baby seat placed on their parents’ lap. All participants viewed the testing material from a distance of 60 to 65 cm.

For infants, the experiment was finished when the trials ended or the infant became fussy and stopped looking at the screen. All participants were monitored during the experiment via a built-in camera, but only infant testing sessions were video-recorded. All adult participants watched the entire stimulus presentation. In the second experiment, adults did not receive any instructions about the task. This way, we wanted to ensure that both infants and adults had the same amount of

prior knowledge about the task specifics in order to robustly measure age related differences in learning strategies within the experiment. After the experiment, adult participants were asked whether they detected any structure in the paradigm and their reports were written down. Participants were debriefed afterwards.

2.2.1. Experiment with adults using instructions

We invited another group of adults for a third experiment in which we gave them an information sheet about the task before they participated in the experiment. Participants were informed that targets (i.e. differently colored bees) would appear on one of the eight designated spots on the circle and a specific sound would accompany each target. They were informed that mostly the targets appeared at the same location for several trials in a row (i.e. expected trials). However, in some trials, the targets appeared on a different location on the circle (i.e. surprise/update trials). When this happened, the color of the targets would also change, so they knew that they should now expect targets to appear at a different location from what they had previously been expecting. They were further informed that, on some of the trials, the targets would appear at random positions on the circle only once. However, it is important to note that participants were not informed about the specific location where they should expect the targets to appear following surprise/no-update trials.

2.3. Data processing

We applied exactly the same protocol to process infant and adult data for all experiments. We used the Tobii IV-T Fixation Filter to define fixations. Custom-made Python scripts (version 3.2) were used to process the fixation data. We defined eight areas of interest (AOI) with 140-pixel radius around each of the eight designated spots on the circle (Fig. 2). We used fixation latency as the dependent measure, which was defined as the latency between the appearance of the bee and the first fixation in this AOI.

We first segmented fixation events for each trial window (i.e. 3000 ms per trial). Minimum fixation duration was set to 100 ms and fixations below this value were disregarded. To ensure that gaze shifts to the target were made from the central fixation image, only trials were included in which a valid fixation was detected on the fixation image first and then on the target AOI during the same trial window. In

order to account for the variability in the infant data, it was necessary to define an individualized familiarization period for each participant. For this, we targeted the first valid surprise/update and surprise/no-update trials for each participant, which included at least one expected trial in between. This window was regarded as the “familiarization phase” and values after this point were taken into account for the main analyses. Participants who provided data for at least one more surprise/update and surprise/no-update trial (separated by expected trials) besides the familiarization phase were included in the final sample (infants: expected trials, $M = 37.62$, $SD = 17.49$; surprise/update trials $M = 5.38$, $SD = 2.68$, surprise/no-update trials, $M = 4.45$, $SD = 2.38$). For adults, this phase almost always corresponded to the initial trials in the task (i.e. approximately the first 8 trials), because adults registered valid fixations almost during the entire task. To remove outliers, we disregarded fixation latencies which were more than double the average value and larger than 3 SDs from the mean (cf. O’Reilly et al., 2013).

3. Results

We had clear predictions about the adults’ performance in the task based on the previous research (cf. O’Reilly et al., 2013) and examined whether infants would already show similar pattern of responses. As shown in Fig. 3, we predicted that participants would be slower performing an eye movement to surprising targets because in these trials the targets appeared at a location that was different from the previous trials, which required re-programming of the eye movement response resulting in increased fixation latencies. If participants learned the structure of the experiment and differentially represented the predictive value of the color cues, we expected them to look at the target locations in the trials following surprising trials (i.e. trials +1, +2, +3 and so forth) significantly faster than in response to surprising trials. The pattern of a steep decrease after a surprising trial for both surprise/update and surprise/no-update trials indicates that in trials +1, a correct prediction about the trial location is generated, which differs for surprise/update and surprise/no-update trials. In other words, if participants failed to update their predictions on surprise/update trials, or conversely, if they did update their predictions on surprise/no-update trials, they would expect the wrong location on trials +1, and thus, have slower fixation latencies on that trial in the condition in which the

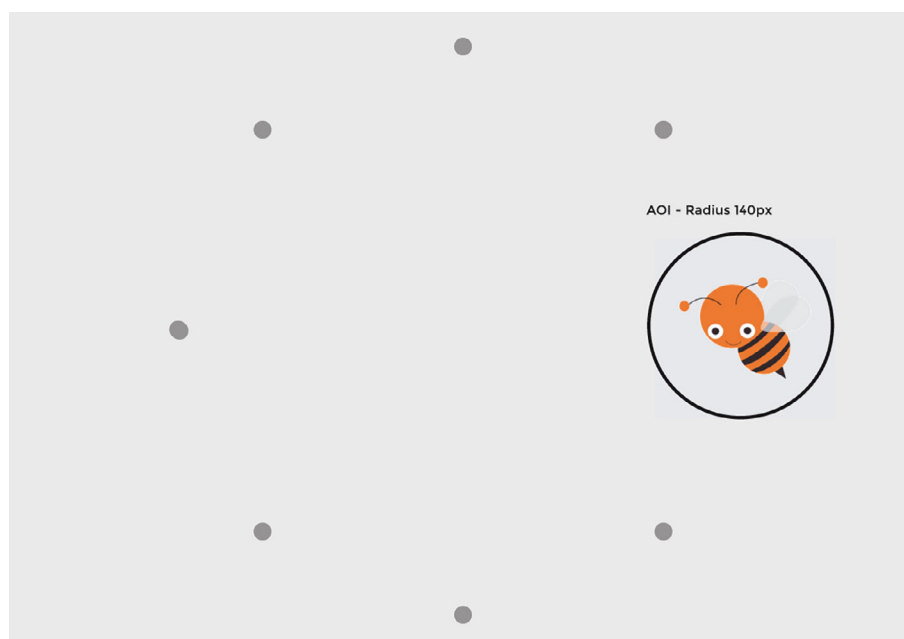


Fig. 2. An example stimulus image that shows the position and size of an area of interest (AOI).




Fig. 3. An illustration of the predicted fixation latencies on surprise/update, surprise/no-update and surrounding trials (i.e. expected trials) if participants differentiated the two types of surprising trials. Trials 0 indicate surprising trials (surprise/update or surprise/no-update). Trials -1 are the trials immediately before, which were expected, whereas trials +1, +2 and +3 are the trials immediately following surprising trials. Please note that if trial 0 is a surprise/update trial, the target stimulus of trial +1 appears at the same location as in trial 0, whereas if trial 0 is a surprise/no-update trial, the target of trial +1 appears at the same location as in trial -1.

error had occurred.

3.1. Experiment-1: Infant experiment

To test our hypotheses we first ran a 2 × 4 repeated measures ANOVA with trial type (surprise/update, surprise/no-update) and trial number (0: surprising trials, trials +1, 2, 3: immediately following trials) as within-subjects factors. This analysis revealed a significant main effect of trial number ($F(3, 63) = 5.41, p = 0.002, \eta^2 = 0.21$) and a marginally significant main effect of trial type ($F(1, 21) = 3.47, p = 0.077, \eta^2 = 0.14$). Data revealed no trial type by trial number interaction ($F(3, 63) = 0.20, p = 0.896, \eta^2 = 0.01$) suggesting that infants' fixation latencies did not differ significantly for the two types of surprising trials over the course of trials, which confirmed our hypothesis (compare Figs. 3 and 4).

We ran follow up tests using the Least Significant Differences method to examine the main effect of trial number. These analyses revealed that infants were significantly faster to look at the target




Fig. 4. Fixation latencies on surprise/update, surprise/no-update and surrounding trials (i.e. expected trials) for 14-month-old infants. Trials 0 indicate surprise/update and surprise/no-update trials themselves. Trials -1 are the trials immediately before surprise/update or surprise/no-update trials, whereas trials +1, +2 and +3 are the trials immediately following surprising trials. Error bars indicate the standard error of the mean.




Fig. 5. Fixation latencies on surprise/update, surprise/no-update and surrounding trials (i.e., expected trials) for uninstructed adults. Trials 0 indicate surprise/update and surprise/no-update trials themselves. Trials -1 are the trials immediately before surprise/update or surprise/no-update trials, whereas trials +1, +2 and +3 are the trials immediately following surprising trials. Error bars indicate the standard error of the mean.

locations following surprising trials (trials 0 vs. trials +1: $md = 46.08, SE = 19.56, p = 0.028$; trials 0 vs. trials +2: $md = 58.42, SE = 17.96, p = 0.004$; trials 0 vs. trials +3: $md = 80.59, SE = 18.35, p = 0.000$).

3.2. Experiment-2: Adult experiment without instructions

We had predicted that both, infants and adults would show a pattern of responses as illustrated in Fig. 3, if they learn the structure of the experiment and dissociate between the different types of unexpected information. The data of the adult participants, however, revealed a different response pattern (compare Figs. 3 and 5). Verbal reports indicated that 13 out of 23 adult participants did not detect any structure in the paradigm and the other 10 participants described diverse complex patterns that were not correct.

To test the response patterns statistically, we used the same analyses as for the infant experiment. Because Mauchly's Test of Sphericity indicated that the assumption of sphericity had been violated ($p < 0.05$), we used Greenhouse-Geisser corrections. A repeated measures ANOVA with trial type (surprise/update, surprise/no-update) and trial number (0: surprising trials, trials +1, 2, 3: immediately following trials) revealed a significant main effect of trial type ($F(1, 22) = 37.08, p < 0.001, \eta^2 = 0.63$) and a significant main effect of trial number ($F(1.45, 31.93) = 15.19, p < 0.001, \eta^2 = 0.41$). Data also revealed a significant trial type by trial number interaction ($F(2.20, 48.29) = 12.46, p < 0.001, \eta^2 = 0.36$). This suggests that adults' fixation latencies differed for the two trial types over the course of trials.

Follow up tests using the Least Significant Differences method for pairwise comparisons of estimated marginal means revealed that adults' responses in surprising trials significantly differed from the immediately following trials (0 vs. trials +1: $md = 32.05, SE = 9.12, p = 0.002$; trials 0 vs. trials +2: $md = 60.44, SE = 15.46, p = 0.001$; trials 0 vs. trials +3: $md = 70.77, SE = 15.59, p < 0.001$). Moreover, they were significantly slower in surprise/no-update trials as compared to the surprise/update trials ($md = 38.04, SE = 6.25, p < 0.001$).

Because data revealed a significant trial type by trial number interaction, we investigated adults' response patterns in different trials separately for each trial type. We ran repeated measures ANOVAs for each trial type with the trial number (0: surprising trials, trials +1, 2, 3: immediately following trials) as the within subjects factor, which were followed up by tests using the Least Significant Differences method for pairwise comparisons of estimated marginal means. These analyses revealed a significant main effect of trial number both in surprise/

update ($F(1.99, 43.85) = 12.69, p < 0.001, \eta^2 = 0.37$) and surprise/no-update trials ($F(1.62, 35.62) = 16.53, p < 0.001, \eta^2 = 0.43$). Similar to infants, adults were significantly faster to fixate the targets in the trials immediately following surprise/update trials (trials 0 vs. +1: $md = 78.33, SE = 17.98, p < 0.001$; trials 0 vs. +2: $md = 78.71, SE = 19.87, p = 0.001$; trials 0 vs. +3: $md = 76.05, SE = 18.31, p < 0.001$).

Whereas adults showed the predicted response pattern for trials following the surprise/update trials, this was not the case for surprise/no-update trials (compare Figs. 3 and 5). Interestingly, unlike infants, adult participants seemed to struggle to learn about the surprise/no-update trials and showed marginally slower fixation latencies in the trials immediately following surprise/no-update trials as compared to the preceding surprise/no-update trials (trials 0 vs. +1: $md = -14.23, SE = 7.13, p = 0.059$). Nevertheless, they became significantly faster to look at the target locations in the following trials (trials 0 vs. +2: $md = 42.17, SE = 13.26, p = 0.004$; trials 0 vs. +3: $md = 65.49, SE = 14.62, p < 0.001$).

3.2.1. Additional analyses for Trials + 1

Our analyses revealed that there was no indication that adults dissociated the two types of surprising trials. Given these unexpected results, we explored the data further to investigate what kind of strategies adults employed during the experiment. We reasoned that if participants had a tendency to update their predictions erroneously in surprise/no-update trials, they would first look at the location of the surprise/no-update bee in the immediately following trials, although the bees did not appear at this location in these trials, which would result in increased fixation latencies. For comparison, we focused on the gaze behavior of both infants and adults during the trials immediately following surprise/no-update trials and systematically investigated the gaze patterns during this time window¹. We first determined the number of infant and adult participants who shifted their gaze at least once to the location of the surprise/no-update trials (i.e. at the location of trials 0) also on the immediately following trials (i.e. trials + 1). Data revealed that among 29 infants, nine infants expected the following target to appear at the location where the surprise/no-update target appeared at least once (range 1 to 3 times). Put differently, on average, only in 5.99% of all trials + 1 following surprise/no-update trials, infants expected the target to appear at the same location as the previous surprise/no-update target location (i.e. at the location of trials 0). These analyses provide further evidence that infants learned the structure of the experiment and adjusted their models accordingly.

The adults, however, showed a different pattern. Among 23 adults, 20 of them expected the target to appear at the same location as the previous surprise/no-update trial, in the trial after, at least once (range 1 to 11 times out of 14). In adults, the average percentage of trials + 1 in which they performed a gaze-shift to the location of the previous surprise/no-update location was 29.17%. These results show that, without instructions, adults had a tendency to update their predictions following surprise/no-update trials approximately 5 times more than the infants did. Together, these results indicate that it was more challenging for adults than infants to differentiate the two types of surprising trials. Instead, they expected the target of the trials following all surprising trials to occur at the same location as the surprising trial itself. Further analyses comparing the results of infants and

¹ To process this data, we defined a second area of interest (AOI) to check if participants gazed at the location of the agent in the preceding trial, in trials following surprise/no-update trials. This additional AOI was active while the bee image was on the screen, in the trial immediately following surprise/no-update trials. We checked the number of trials per participant when they gazed at this second AOI. To calculate the percentage values, we divided the number of trials + 1, in which participants gazed at this second AOI, by the total number of trials + 1 for each participant, multiplied this value with 100 and took an average of all percentage values.

uninstructed adults are reported in [Supplementary Material](#).

3.3. Experiment-3: Adult experiment with instructions

Because adults showed different patterns than predicted by our hypothesis, we tested another group of adults who were provided with instructions about the task. We expected adults to perform the task as predicted (i.e. they dissociate the two types of unexpected trials) when the uncertainty about the task demands were reduced (cf. O'Reilly et al., 2013). For statistical testing, we applied the same analysis protocol as reported earlier. Mauchly's Test of Sphericity indicated that the assumption of sphericity had been violated ($p < 0.05$); therefore, we used Greenhouse-Geisser corrections. A repeated measures ANOVA with trial type (surprise/update, surprise/no-update) and trial number (0: surprising trials, trials + 1, 2, 3: immediately following trials) indicated a significant main effect of trial number ($F(1.75, 43.64) = 45.64, p < 0.001, \eta^2 = 0.65$) and a significant main effect of trial type ($F(1, 25) = 24.13, p < 0.001, \eta^2 = 0.49$). Follow-up tests showed that adults were significantly faster in the trials immediately following surprising trials (trials 0 vs. trials + 1: $md = 124.24, SE = 17.69, p < 0.001$; trials 0 vs. trials + 2: $md = 158.35, SE = 20.25, p < 0.001$; trials 0 vs. trials + 3: $md = 145.78, SE = 19.06, p < 0.001$). Furthermore, they were significantly slower in surprise/no-update trials as compared to the surprise/update trials ($md = 28.75, SE = 5.85, p < 0.001$).

Data also revealed a significant trial type by trial number interaction in this group ($F(1.92, 48.06) = 9.99, p < 0.001, \eta^2 = 0.29$). As illustrated in Fig. 6, this interaction effect was driven by the slower fixation latencies following surprise/no-update trials. This might be because participants were specifically informed about where the targets appear following surprise/update trials (i.e. expected trials), which was not the case for the surprise/no-update trials. Participants were only informed that following surprise/no-update trials, the targets would not appear where they appeared in the preceding surprise/no-update trial.

To have a closer look at the response patterns in different trial types, we ran separate repeated measures ANOVAs for each trial type with trial number (0: surprising trials, trials + 1, 2, 3: immediately following trials) as the within subjects factor, which we followed up with tests using the Least Significant Differences method for pairwise comparisons of estimated marginal means. Data revealed a significant main effect of trial number both in surprise/update ($F(1.94, 48.53) = 46.48, p < 0.001, \eta^2 = 0.65$) and surprise/no-update trials ($F(3, 75) = 29.25, p < 0.001, \eta^2 = 0.54$). Fixation latencies of adults significantly decreased in the trials immediately following surprise/update trials (trials 0 vs. trials + 1: $md = 163.51, SE = 20.61, p < 0.001$; trials 0 vs. trials + 2: $md = 170.36, SE = 21.35, p < 0.001$; trials 0 vs. trials + 3: $md = 143.33, SE = 19.47, p < 0.001$). Moreover, participants were significantly faster to look at the target locations also in the trials immediately following surprise/no-update trials as compared to the surprise/no-update trials (trials 0 vs. trials + 1: $md = 84.96, SE = 19.94, p < 0.001$; trials 0 vs. trials + 2: $md = 146.34, SE = 20.94, p < 0.001$; trials 0 vs. trials + 3: $md = 148.23, SE = 20.03, p < 0.001$).

4. Discussion

In three experiments, we examined whether infants and adults represent the structure of a changing experimental environment. Infants learned the predictive value of different color cues (i.e. update vs. no-update) and adjusted their models when required. Interestingly, unlike infants, especially when not guided by instructions, adults seemed to have a tendency to update their models throughout the experiment. When adults received instructions about the task, their performance was slightly improved.

Although infants and adults who were not instructed had exactly the same prior knowledge about the task specifics, we observed clear




Fig. 6. Fixation latencies on surprise/update, surprise/no-update and surrounding trials (i.e., expected trials) for instructed adults. Trials 0 indicate surprise/update and surprise/no-update trials themselves. Trials -1 are the trials immediately before surprise/update or surprise/no-update trials, whereas trials +1, +2 and +3 are the trials immediately following surprising trials. Error bars indicate the standard error of the mean.

differences in their responses. We argue that this finding might be explained by infants being more open to accommodate current information than adults were. It could be that infants are more ready to learn from evidence and adjust their beliefs flexibly, due to having less precise beliefs about the parameters of the environment (Courville et al., 2006). Adults, on the other hand, might have adhered to one strategy to reduce the uncertainty and ignored some of the available evidence. This argument is further supported by the additional analyses showing that after surprise/no-update trials, uninstructed adults initially expected the target to appear at the surprise/no-update location, just like after surprise/update trials. They then looked at the correct target location, resulting in increased fixation latencies in trials immediately following surprise/no-update trials. This response pattern shows that adults had a tendency to update their models following all unexpected items, despite the evidence suggesting a differentiation between unexpected cues.

Relatedly, one might claim that because of poorer memory skills, infants were not able to represent prior information and their responses were driven solely by the current data (i.e. simply looking at a spot where something appears). Therefore, one could argue that a mechanism based on *forgetting* previous information might have helped infants to succeed on this task. Although neural and modeling work would be useful to explore this argument further, it is insufficient to account for infants' pattern of responses. If infants' looking responses were only driven by current observations, we would have observed an even response pattern across all trials. First, infants would be equally fast to look at the target locations regardless of whether the target appears at an unexpected location or not. That is, they would not show increased fixation latencies in surprising trials, if they did not use prior information about where the targets appear. Second, if infants did not learn that targets appear at different locations based on whether the previous trial is an update or no-update trial, they would not be fast again to look at the correct target locations in the trials immediately following these trials (see Fig. 4). Therefore, it is unlikely that infants did not encode any previous information to generate predictive models and their responses were driven purely by the current input.

The differences between infants and adults show fascinating similarity to findings suggesting that young children use evidence to learn novel causal links better than adults do (Gopnik, Griffiths, & Lucas, 2015; Lucas, Bridgers, Griffiths, & Gopnik, 2014). For example, in a study by Lucas et al. (2014), adults and 4-year-old children had to infer how a machine played music when some objects were placed on top of it. In some cases, a single object activated the machine (i.e. individual principle) whereas in other cases it was a combination of objects (i.e.

combination principle) activating the machine. Children learned both principles to make the machine play music and selected the right principle when they had to infer which principle should work. However, adults tended to stick to the individual principle that they learned earlier and placed single objects to activate the machine, although current evidence was against this principle. Similarly, in the current paradigm, particularly without instructions, adult participants seemed to adhere to one explanation (i.e. the unexpected appearance of the bees means an update of location), although evidence suggested a differentiation between different forms of surprising information. Because adults exploited one strategy (i.e. update), they were less open to exploring alternative explanations, which hindered them from performing the task successfully.

One alternative explanation could be that although the number of surprise/update and surprise/no-update trials was matched, overall, the bee stimulus used on surprise/update trials was seen more frequently than the surprise/no-update bees, as the same bee appeared at the same spot as “expected” bees in the subsequent trials. Accordingly, it might be that participants, particularly infants, might have failed to learn about the rare surprise/no-update items or simply ignored them as distractors. One way to address this limitation could have been to use a fourth color of the bees to signal surprise/update trials, which could have made the task less understandable, especially for the infants.

Despite this limitation, our results suggest that infants did treat surprise/update trials differently from surprise/no-update trials, which is observed in decreased fixation latencies in the trials following the surprising trials. If infants did not learn what the surprise/update colors represented, they would not be fast to look at the correct bee locations in the trials after the surprise/update trials. Similarly, if infants did not learn that surprise/no-update bees were *not* predictive of future bees, they would not be equally fast to look at the correct bee locations in the trials following the surprise/no-update trials. This is especially crucial given that in these trials both the color and the location of the bees were different from the preceding (surprise/no-update) trials incurring a larger reaction time cost than the trials after the surprise/update trials. Therefore, it would be reasonable to assume that infants learned that the surprise/update and surprise/no-update bees have different predictive value with regard to which stimuli they should expect on the following and subsequent trials.

Another alternative explanation could be that, due to having complex priors, adult participants failed to use the information value of the simple perceptual cues such as color to perform the task as expected. In other words, the complex prior knowledge adults brought to the experiment might have dominated their predictions, which resulted in overlooking the perceptual cues. Relatedly, it could be argued that, especially in the absence of instructions, adults had to detect the structure (e.g., an update model) and the parameters (e.g., the target will now appear where it just appeared) of the current task to make predictions accordingly (Braun, Mehring, & Wolpert, 2010). For adults, detecting the structure of the environment (and the parameters thereof) might have been a challenging task because they had to narrow down a large hypothesis space to figure out the general structure and the parameters of the environment². For infants, the task might have been easier, as there is ample evidence in the developmental literature showing that infants detect even complex statistical regularities (Kirkham, Slemmer, & Johnson, 2002; Ruffman, Taumoepeau, & Perkins, 2012; Saffran, Aslin, & Newport, 1996) and younger learners are better than the older ones and adults (Kuhl, 2004; Lucas et al., 2014). In summary, it might be that having more prior knowledge prevented the adults to perform the task as expected.

It should be noted that results obtained from adult participants in

² This argument is supported by the verbal reports from uninstructed adults. Adults who reported structure detection tried to describe diverse complex patterns that were incorrect.

the current study show differences to the findings of the 2013 study by O'Reilly and colleagues. Several methodological differences regarding the number of trials, the stimuli and duration of the presentation might have led to the discrepant results, which make it difficult to compare the studies. First, in the O'Reilly study, participants completed a 300-trials session for the eye tracking study whereas, in the current study, due to infants' limited attention span, it was not possible to include many trials in the experiment. Moreover, O'Reilly and colleagues used dots in their task that appeared at a faster rate and stayed on the screen relatively short (i.e. 350 ms). One could argue that, in the current study, because the presentation of the images were slow (e.g. the bees stayed on the screen for 1.5 s), it was more likely that adults had more time to think of more complex explanations, which might have resulted in overlooking simple perceptual cues. In addition, in the O'Reilly study, the update dots were chromatically colored, whereas the one-off dots were gray, which might have made the difference between the two types of surprising trials more salient, as it was clear for the participants that colored dots signaled updating whereas gray dots did not. We used differently colored bees to keep infants' attention directed to the screen, which might have made the distinction between the two types of surprising trials less salient. Taken together, although the 2013 study by O'Reilly and colleagues inspired the current work, there were several differences in the methods, which make it difficult to discuss replicability.

5. Conclusions

Here, we demonstrated that infants and adults generate internal models that represent the statistics of a dynamic environment. Infants learned the predictive value of different cues and adjusted their models only when the cues signaled a change in the structure. Unlike infants, especially when not guided by instructions, adults ignored the evidence and had a tendency to update their models each time they observed a change in the structure. We argue that infants are more open to learning from evidence and update their beliefs flexibly, which is an advantageous learning strategy to represent dynamic environments. This strategy enables developing systems to adjust to changes easily, a capacity that is key to adapting to novel environments.

Acknowledgements

We would like to thank the infants, families and adults who participated in the study, and we are grateful to Angela Khadar for organizing the recruitment of participants and Anne Gerrits for her assistance in data collection for the third experiment. We thank the European Union Seventh Framework Programme (grant number 289404) for their funding.

Appendix A. Supplementary material

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2019.02.004>.

References

- Braun, D. A., Mehring, C., & Wolpert, D. M. (2010). Structure learning in action. *Behavioural Brain Research*, 206(2), 157–165. <https://doi.org/10.1016/j.bbr.2009.08.031>.
- Courville, A. C., Daw, N. D., & Touretzky, D. S. (2006). Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences*, 10(7), 294–300. <https://doi.org/10.1016/j.tics.2006.05.004>.
- Draganova, R., Eswaran, H., Murphy, P., Huotilainen, M., Lowery, C., & Preissl, H. (2005). Sound frequency change detection in fetuses and newborns, a magnetoencephalographic study. *Neuroimage*, 28(2), 354–361. <https://doi.org/10.1016/j.neuroimage.2005.06.011>.
- Emberson, L. L., Richards, J. E., & Aslin, R. N. (2015). Top-down modulation in the infant brain: Learning-induced expectations rapidly affect the sensory cortex at 6 months. *Proceedings of the National Academy of Sciences*, 112(31), 9585–9590. <https://doi.org/10.1073/pnas.1510343112>.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>.
- Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When younger learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*, 24(2), 87–92. <https://doi.org/10.1177/0963721414556653>.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35–B42. [https://doi.org/10.1016/S0010-0277\(02\)00004-5](https://doi.org/10.1016/S0010-0277(02)00004-5).
- Kouider, S., Long, B., Le Stanc, L., Charron, S., Fievet, A. C., Barbosa, L. S., & Gelskov, S. V. (2015). Neural dynamics of prediction and surprise in infants. *Nature Communications*, 6. <https://doi.org/10.1038/ncomms9537>.
- Kuhl, P. K. (2004). Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience*, 5(11), 831–843. <https://doi.org/10.1038/nrn1533>.
- Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, 131(2), 284–299. <https://doi.org/10.1016/j.cognition.2013.12.010>.
- O'Reilly, J. X. (2013). Making predictions in a changing world-inference, uncertainty, and learning. *Frontiers in Neuroscience*, 7(105), 1–10. <https://doi.org/10.3389/fnins.2013.00105>.
- O'Reilly, J. X., Schuffelgen, U., Cuell, S. F., Behrens, T. E., Mars, R. B., & Rushworth, M. F. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 110(38), E3660–E3669. <https://doi.org/10.1073/pnas.1305373110>.
- Partanen, E., Kujala, T., Näätänen, R., Liitola, A., Sambeth, A., & Huotilainen, M. (2013). Learning-induced neural plasticity of speech processing before birth. *Proceedings of the National Academy of Sciences*, 110(37), 15145–15150. <https://doi.org/10.1073/pnas.1302159110>.
- Payzan-LeNestour, E., & Bossaerts, P. (2011). Risk, unexpected uncertainty, and estimation uncertainty: Bayesian learning in unstable settings. *PLoS Computational Biology*, 7(1), e1001048. <https://doi.org/10.1371/journal.pcbi.1001048>.
- Picard, F., & Friston, K. (2014). Predictions, perception, and a sense of self. *Neurology*, 83(12), 1112–1118. [10.1212/WNL.0000000000000798](https://doi.org/10.1212/WNL.0000000000000798).
- Ruffman, T., Taumoepeau, M., & Perkins, C. (2012). Statistical learning as a basis for social understanding in children. *British Journal of Developmental Psychology*, 30(1), 87–104. <https://doi.org/10.1111/j.2044-835x.2011.02045.x>.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926–1928. <https://doi.org/10.1126/science.274.5294.1926>.
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology*, 69, 181–213. <https://doi.org/10.1146/annurev-psych-122216-011805>.
- Sambeth, A., Pakarinen, S., Ruohio, K., Fellman, V., van Zuijen, T. L., & Huotilainen, M. (2009). Change detection in newborns using a multiple deviant paradigm: A study using magnetoencephalography. *Clinical Neurophysiology*, 120(3), 530–538. <https://doi.org/10.1016/j.clinph.2008.12.033>.
- Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. *Science*, 348(6230), 91–94. <https://doi.org/10.1126/science.aaa3799>.