Taylor & Francis
Taylor & Francis Group

Connection Science

ISSN: 0954-0091 (Print) 1360-0494 (Online) Journal homepage: http://www.tandfonline.com/loi/ccos20

How infants' reaches reveal principles of sensorimotor decision making

Evelina Dineva & Gregor Schöner

To cite this article: Evelina Dineva & Gregor Schöner (2018) How infants' reaches reveal principles of sensorimotor decision making, Connection Science, 30:1, 53-80, DOI: 10.1080/09540091.2017.1405382

To link to this article: https://doi.org/10.1080/09540091.2017.1405382

Published online: 30 Jan 2018.

CONNECTION SCIENCE, 2018
VOL. 30, NO. 1, 53–80
https://doi.org/10.1080/09540091.2017.1405382

How infants' reaches reveal principles of sensorimotor decision making

Evelina Dineva and Gregor Schöner

Institute for Neural Computation, Ruhr-Universität Bochum, Bochum, Germany

ABSTRACT

In Piaget's classical *A*-not-*B*-task, infants repeatedly make a sensorimotor decision to reach to one of two cued targets. Perseverative errors are induced by switching the cue from *A* to *B*, while spontaneous errors are unsolicited reaches to *B* when only *A* is cued. We argue that theoretical accounts of sensorimotor decision-making fail to address how motor decisions leave a memory trace that may impact future sensorimotor decisions. Instead, in extant neural models, perseveration is caused solely by the history of stimulation. We present a neural dynamic model of sensorimotor decision-making within the framework of Dynamic Field Theory, in which a dynamic instability amplifies fluctuations in neural activation into macroscopic, stable neural activation states that leave memory traces. The model predicts perseveration, but also a tendency to repeat spontaneous errors. To test the account, we pool data from several *A*-not-*B* experiments. A conditional probabilities analysis accounts quantitatively how motor decisions depend on the history of reaching. The results provide evidence for the interdependence among subsequent reaching decisions that is explained by the model, showing that by amplifying small differences in activation and affecting learning, decisions have consequences beyond the individual behavioural act.

ARTICLE HISTORY
Received 29 March 2017
Accepted 18 September 2017

KEYWORDS
Decision making; neural dynamics; dynamic field theory; infant development spontaneous errors

1. Introduction

Consider the classical *A*-not-*B* effect discovered by Piaget (1954) and replicated numerous times (see Wellman, Cross, & Bartsch, 1987, for a review). An infant is presented with a box with two hiding locations, *A* and *B*. While the infant is watching, the experimenter hides a toy at the "*A*" location and, after a brief delay, pushes the box close enough for the infant to reach. The infant typically reaches to the *A* location. After a few of such *A* trials, the experimenter switches and hides the toy at the *B* location. Young infants of about 7 months typically perseverate and reach to *A*, making the *A*-not-*B* error, while older infants of about 11 months reliably reach to *B*.

CONTACT Gregor Schöner gregor.schoener@ini.rub.de Institute for Neural Computation, Ruhr-Universität Bochum, 44870 Bochum, Germany

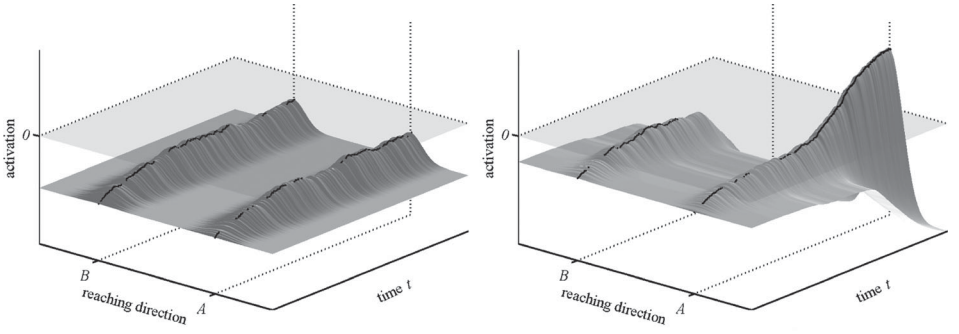

Figure 1. Neural activation fields representing reaching direction plotted as functions of time. A field with relatively week inputs (left panel) is compared to a field with stronger inputs (right panel); the latter one generates a self-stabilised peak of activation. At any moment in time, the location of maximal activation is marked by a black dot for either field.

At the core of this task is a sensorimotor decision to generate one of two reaches. This is clear in the toy-less version of the *A*-not-*B* task (Smith, Thelen, Titzer, & Mclin, 1999), in which the experimenter attracts attention to the *A*, and later to the *B* location by waving the lid that covers the corresponding hiding location. As in the original version of the task, infants' reach for the lid displays perseveration on *B* trials. Neural accounts of such sensorimotor decisions entail neural activation patterns, in which neurons voting for either reach have different levels of activation and some kind of threshold mechanism decides who is the winner (Gold & Shadlen, 2007). Figure 1 illustrates a neural account of the decision between *A* and *B* that is framed within Dynamic Field Theory (DFT, Schöner, Spencer, & the DFT research group, 2016). In that account, neural populations representing the movement plan (Georgopoulos, 1990) are modelled by an activation field defined over reaching direction (Erlhagen, Bastian, Jancke, Riehle, & Schöner, 1999). In one conception, visualised on the left, the decision is taken by picking the reaching direction with largest activation. In that left panel of the figure, two locations have the same mean level of activation, but activation fluctuates, so when the decision is "read-out" at some fixed time (e.g. at the end of the delay), either choice is equally possible. This is how the original DFT model of perseverative reaching accounted for chance level performance (Thelen, Schöner, Scheier, & Smith, 2001). In the nervous system, however, motor actions are not based on a one-time reading of a neural representation. It takes time to initiate a motor action, and during that time, the motor system is continuously linked to sensory inputs, a feature known as "online updating" (Georgopoulos, Lurito, Petrides, Schwartz, & Massey, 1989; Goodale, Pélisson, & Prablanc, 1986). Continuously "reading out" the most activated location during a time interval would lead to wildly fluctuating choices, however, as the most activated location may switch randomly from moment to moment. In fact, a robotic implementation of original DFT account that we attempted a few years ago ran into exactly this problem: The robot's movement was erratic even within a single trial as the control signal kept switching randomly between the two targets.

The right panel of Figure 1 illustrates an alternative conception. Here, microscopic random fluctuations are transformed into a macroscopic difference between different

outcomes that represents the decision. This is based on neural interaction within the field. A fluctuation around the *A* location is amplified into a macroscopic peak by excitatory interaction among neighbouring field sites. As activation grows around the *A* location, it begins to suppress activation at other locations through inhibitory interaction. The activation peak that results is a stable state that persists in time, so that within a trial the sensorimotor decision is stable. Continuously coupling this activation peak to a motor system generates stable movement behaviour. On another trial, however, the peak may arise at the other site if a chance fluctuation raises the level of activation at *B* rather than *A*. So this is still an account for chance level performance, but now consistent with the stability demands of real motor behaviour. In fact, the self-stabilised peak can be directly coupled into a behavioural dynamics that moves an effector (Bicho, Mallet, & Schöner, 2000). The account is "autonomous" in the sense that it does not rely on an additional cognitive process or agent that (i) detects that the time for making the decision has come, (ii) "reads" the most activated choice, and (iii) translates that reading into a motor command.

This seemingly subtle issue has important consequences. The very notion of perseveration is that past reaching decisions influence present reaching decisions. In neural terms, each reaching decision must leave a trace capable of exerting this influence. Such a memory trace is neurally plausible only if the difference in activation between the selected and the alternate reach is macroscopic and persists in time beyond a single instant of "read-out". In fact, in the original DFT model, reaching decisions did not leave a memory trace and perseveration was not a reflection of past reaching decisions, but a reflection of past patterns of stimulation (Thelen et al., 2001). The same is true of the model of Munakata (1998) within the framework of parallel distributed processing (PDP) (Rumelhart, McClelland, & The PDP Research Group, 1986).

In this paper, we make explicit the requisite neural mechanisms for making stable sensorimotor decisions that leave a memory trace. The key insight is that young infants make the sensorimotor decision only at the end of the delay, when the cue to *A* or to *B* is no longer the dominant source of specification. This makes the reaching decision in the *A*-not-*B*-paradigm sensitive to the different sources of specification. As a result of this insight, perseverative reaching in infants is recognised as a window into understanding sensorimotor decision making that complements the paradigms of speeded decision making that are prevalent in the adult and the neurophysiological literature.

To test the new account we pool experimental data on infant perseverative reaching from a set of studies. This enables us to estimate of a set of conditional probabilities that are predicted by the model. A salient prediction involves "spontaneous errors", in which infants reach to *B* on an *A* trial. We will show that making a spontaneous error increases the probability of making it again and reduces the probability of perseveration.

2. Neural dynamic account for stable sensorimotor decisions

We build on the previous DFT model of perseverative reaching (Thelen et al., 2001), in which the sensorimotor decision is based on a neural activation field defined over reaching direction. The DFT framework emphasises that dynamic stability is critical when sensorimotor or even cognitive processes are linked to the sensory and motor surfaces (Schöner et al., 2016).

Ironically, in the original model (Thelen et al., 2001), this conceptual commitment of DFT was violated. The model looked only at the activation patterns up to the end of the modelled delay period and then simply "read-out" the location of maximal activation along the field dimension, a process that lies outside the conceptual framework of DFT. The model failed to account for the consequences of decisions, an error detected only years after the reception of the original model (Schöner & Dineva, 2007).

Munakata and colleagues modelled perseverative reaching (Munakata, 1998; Munakata, McClelland, Johnson, & Siegler, 1997) within the language of PDP (Rumelhart et al., 1986). The sensorimotor decision is based on the activation level of two nodes that stand for the two possible choices. The decision is "read-out" from the model at the end of the delay by translating the relative difference in activation of the two output neurons into a probability for making either of the two choices Munakata (1998, footnote 5 on page 168). So in this account too, the decision itself is not reflected in the neural activation pattern and does not, therefore, influence the learning process.

The exclusive focus on the processing up to a decision time is, in fact, a typical feature of models of decision making. In a classical diffusion model of a two-choice task (Ratcliff, Van Zandt, & McKoon, 1999), for instance, the decision is determined by which of two boundaries representing the two choices is first reached. This conception has been translated into neural language by treating the boundaries as thresholds for the firing rates of two neurons that represent the different choices (Gold & Shadlen, 2007). How motor decisions are acted out and how they may be reflected in memory traces, and thus, how past decisions have consequences for future decisions, have been secondary topics in this line of thinking.

2.1. The DFT model

To account for stable sensorimotor decision making we repair the conceptual mistake of the previous DFT model (Thelen et al., 2001). Here we provide a conceptual overview of the new variant of the model, emphasising the different dynamic regimes of the field at different points within the task. We leave the mathematical details to an appendix (Appendix 1).

In the DFT account, the process of making a sensorimotor decision is modelled by the evolution in time of an activation field defined over the reaching direction (Figure 1). A localised peak of positive activation (on the right of the right panel) represents a motor plan to reach in the direction specified by the peak's location. Such a peak is a stable state, an attractor, of the neural dynamics that generates the activation time courses. It is stabilised against decay by local excitatory interaction within the peak. Activation within the peak inhibits activation at all other field sites through inhibitory interaction. This stabilises the peak against spreading out limitlessly (Amari, 1977; Erlhagen & Schöner, 2002).

How do such self-stabilised peaks of activation arise? Without a peak, an activation field is in a subthreshold pattern of activation (shown in the left panel of Figure 1). The subthreshold pattern reflects weak inputs, in this case, two localised inputs around the reaching directions *A* and *B*. Peaks arise when and where activation levels first pass through the threshold of a sigmoidal function that engages neuronal interaction. Activation may be pushed through the threshold by localised input or by global input, a homogeneous

boost, across the entire field. In the right panel of Figure 1, such a boost has lifted the entire field to higher levels of activation. This has made the subthreshold activation pattern unstable. When positive activation arises by fluctuation, local excitatory interaction pulls activation up into a self-stabilised peak while at the same time inhibitory interaction suppresses activation elsewhere.

In DFT, only self-stabilised peaks of activation represent decisions that affect downstream processes so that ultimately movement may be generated (Schöner et al., 2016). Moreover, only such self-stabilised peaks leave a memory trace. Both is due to the principle that only sufficiently activated neural populations impact on any other neural process, a principle formalised mathematically through the sigmoidal threshold function.

In the *A*-not-*B* paradigm, three sources of input were identified (Thelen et al., 2001). The two hiding locations, visibly marked by lids, induce activation at the two corresponding reaching directions ("task input"). When attention is attracted to either location (the "cue" given by waving the lid or by hiding the toy), "specific input" induces activation at the corresponding reaching direction. The memory trace of past reaches induces activation at the direction of those past reaches. In the previous DFT model (Thelen et al., 2001), only the specific input is sufficiently strong to induce a self-stabilised peak. It is that peak that leaves a memory trace. In the model of young infants who perseverate, that peak decays during the delay so that the field is in a subthreshold pattern of activation when the reach is initiated. The reaching direction is selected by "reading out" the location with maximal activation level at that moment in time. This read-out itself does not leave a memory trace. So the memory trace does not reflect the eventual reaching decision, but the cue! Moreover, as argued above, performing a reach based on the location of maximal activation at a single moment in time does not conform to the stability requirement of human movement generation (Erlhagen & Schöner, 2002).

The model that we propose in this article corrects this mistake by recognising that the end of the delay is signalled to the infant through a meaningful perceptual change when the box is pushed into the infant's reaching range. Experiments reliably contain an event of this kind that transitions from the delay period, when the infant is unable to reach, to the post-delay period, when the infant is able to reach. This perceptual change is modelled by a boost, a homogeneous input that pushes the field through the detection instability (final "boost" phase, shown in the top two panels of Figure 2). A self-stabilised peak forms at one of the two reaching directions (middle panel). This peak represents the reaching decision, and it is this peak that leaves a memory trace (bottom panel). Where the peak is induced depends on the interplay of the different sources of input as no single source of input is dominant at the end of the delay.

That the reaching decision was not made by the activation field, but by a read-out procedure was a conceptual mistake. This is highlighted by another weakness of the earlier account: It does not distinguish between trials on which infants reach to a location and trials on which infants do not reach to either location ("non-reaches"). Such non-reaches are reported in experiments, although they are typically not analysed. By definition, a read-out procedure always specifies a reach. In the PDP model (Munakata, 1998) the reaching decision is similarly not made within the network. It is likewise made by interpreting activation levels from outside the model, translating relative activation levels into relative probabilities of either reach. That model does not, therefore, account for non-reaches either (see Mareschal, 1998, for a discussion of this point).

2.2. Accounting for sensorimotor decisions

How does the new DFT model account for the sensorimotor decision of the reaching infant? There are two aspects to that decision, initiating a reach (vs. not reaching) and selecting one of the possible reaching targets. In the *A*-not-*B* paradigm, this decision occurs at the end of the delay. In the model, this is the moment when a boost pushes the level of the entire activation field up towards threshold (top panel of Figure 2). The detection instability will generate a self-excited peak, which represents the decision to reach (if that does not

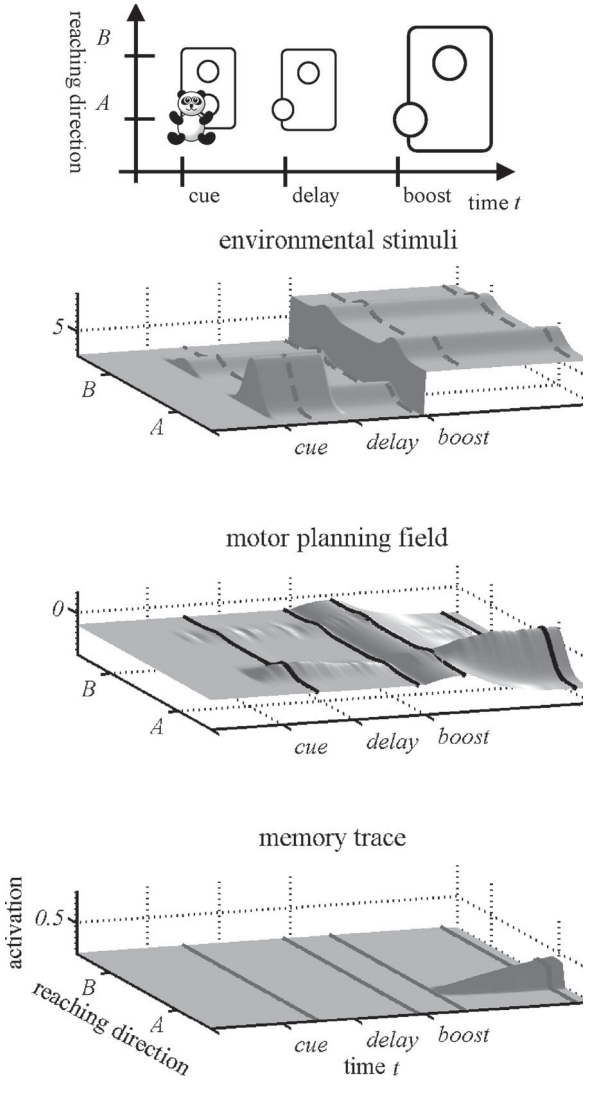

Figure 2. On top, the structure of the first *A* trial with cue, delay and boost phases is illustrated. In the DFT model, this translates into the temporal pattern of perceptual inputs into the field (second from top), which leads to the evolution in time of the field's activation pattern (third from top) and ultimately to the evolution of a memory trace (bottom). The aligned (dashed, solid, and grey) lines indicate concurrent snapshots of activation levels at representative points in time.

happen, we have a non-reach). Selection occurs, because only one peak may form over one of the two regions that receive localised input.

Which of the possible reaches is selected depends on the activation pattern across the field dimension when the boost drives the system through the detection instability. The activation pattern at that time depends on the different inputs to the activation field, as well as on noise. Let us briefly look at each factor impacting on that activation pattern.

Figure 2 shows a case, in which the reaching decision "follows the cue" on an *A* trial. Activation induced around the *A* location by the cue decays during the delay, but may still give the activation level at *A* a slight edge over activation at *B* when the boost arrives. As the boost pushes the field through the threshold, activation at *A* first becomes suprathreshold, beginning to self-excite and to suppress activation at *B*. Varying the cue strength affects how strongly the decision is biased towards the cue. Preliminary results from the model were used to account for experimental data that manipulated cue strength (Clearfield, Dineva, Smith, Diedrich, & Thelen, 2009).

Another reason why the decision may be biased towards *A* could be that the task input is asymmetrical, favouring *A*. This is, in fact, a good description of what happens in many versions of the *A*-not-*B* paradigm, in which experimenters use "training trials" (Figure 3): Over the first few *A* trials, the lids over the *A* location is placed closer to the infant! In the simulations shown in Figure 2, task input was, in fact, asymmetrical in that way, and that

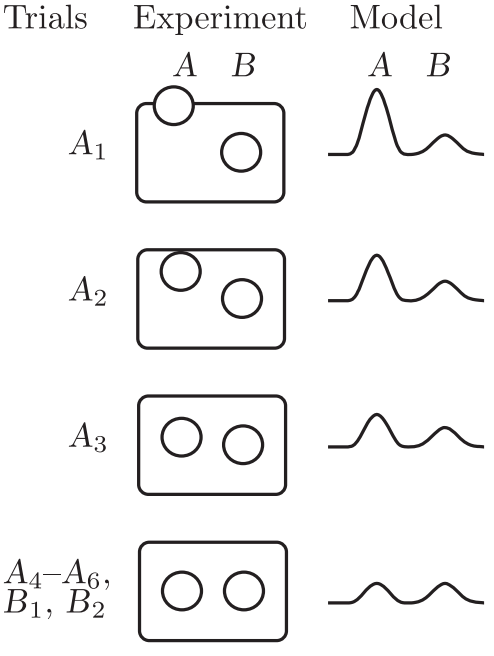

Figure 3. The training procedure used in *A*-not-*B* tasks is shown in the left column: On the first three *A* trials, lid at the *A* location is left closer to the infant after the cue is given until the box is pushed towards the infant. This distance is decreased across the first three *A* trials so that on A_4, the lid at *A* is at the same distance from the infant as the lid at *B*. How the asymmetrical distances of the two lids are simulated by the model is shown in the right column: On the first three *A* trials, task input is stronger at *A* than at *B*. That advantage of *A* is reduced over these first three *A* trials until task input is symmetrical on A_4.

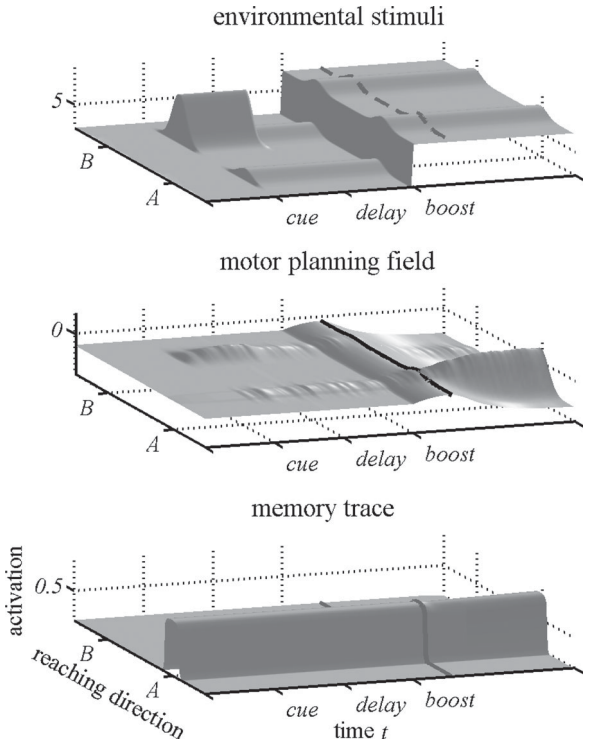

Figure 4. DFT simulation of the first *B* trial (same conventions as in Figure 2).

contributed to the higher activation level at the *A* location, supporting "correct" reaching to *A*. Our quantitative assessment below will exhibit signatures of this dependence in detail.

Finally, the memory trace is a source of input that affects activation levels when the decision is made. Figure 4 shows a simulation of a *B* trial. The cue is given at *B*, but when the boost is given at the end of the delay, input from the memory trace at *A* produces a higher level of activation at *A* than the activation level that remains from the cue at *B*. The self-stabilised peak emerges at *A*, so that the model makes the *A*-not-*B* error.

Neural activation is inherently noisy. In DFT, this fact is accounted for by stochastic inputs that induce fluctuating activation levels across the field. The previous DFT model (Thelen et al., 2001) used only the most generic noise model, independent Gaussian white noise providing input to each field location. This is sufficient to induce stochasticity in the read-out procedure of that earlier model, although the independent noise sources tend to be averaged out by the neural interaction within the field. The present model accounts for the fact that stochasticity originates from the inputs to the activation field (see Appendix 1). This implies, that input noise is spatially correlated and induces stochasticity in the macroscopic activation peaks that emerge from the detection instability. The sensorimotor decisions are, therefore, potentially stochastic in nature. When the differences between activation levels near *A* and *B* are small, stochasticity is observable as a non-zero probability of making the non-dominant decision. Figure 5 shows such a "spontaneous error", a decision for *B* on the fourth *A* trial. On this fourth *A* trial, task input is symmetrical as the "training" phase has ended. Despite differences in input from cue and memory trace, activation levels at *A*

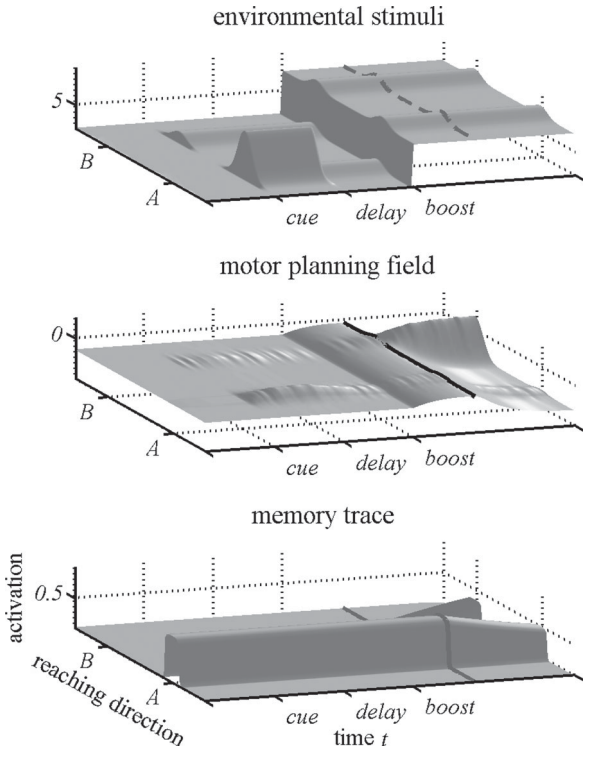

Figure 5. Simulation of trial A_4 of the young infant model with a spontaneous error (same conventions as in Figure 2).

and *B* are quite similar, so that an occasional stochastic fluctuation is sufficient to bring up activation around *B* enough to begin to self-excite and suppress activation near *A*.

Figure 6 illustrates how each sensorimotor decision has consequences for future decisions. Because it is the decision, represented by the self-stabilised peak, that leaves a memory trace, not the input pattern, different stochastic outcomes lead to different memory traces. Two complete runs through the *A*-not-*B* paradigm are shown for the model, six *A* trials followed by two *B* trials. Within each trial, the cue is followed by a delay, which ends with a boost, after which the decision emerges. At the end of each trial, the field is reset to a negative resting level before the next trial begins. The difference between the two runs is only chance, a different outcome of the stochastic decisions. On top, a "regular" pattern of perseverative reaching emerges: On each *A* trial, the model "follows the cue", that is, generates a peak at the *A* location after the delay. The memory trace builds consistently at *A* from each such decision, so that on the *B* trials, activation levels are strongly biased towards *A* at the end of the delay. As a result, on both *B* trials, the model generates peaks at *A*, making the *A*-not-*B* error twice in a row.

On bottom, an "irregular" pattern of reaching emerges. All parameters and stimulus conditions are the exact same as on top. On the fourth *A* trial, however, through a fluctuation, a spontaneous error occurs with a peak building at *B*, rather than at the cued *A* location, at the end of the delay. This different stochastic outcome of the sensorimotor decision leaves a memory trace at *B*, while the trace at *A* decays. In the shown simulation run, a spontaneous

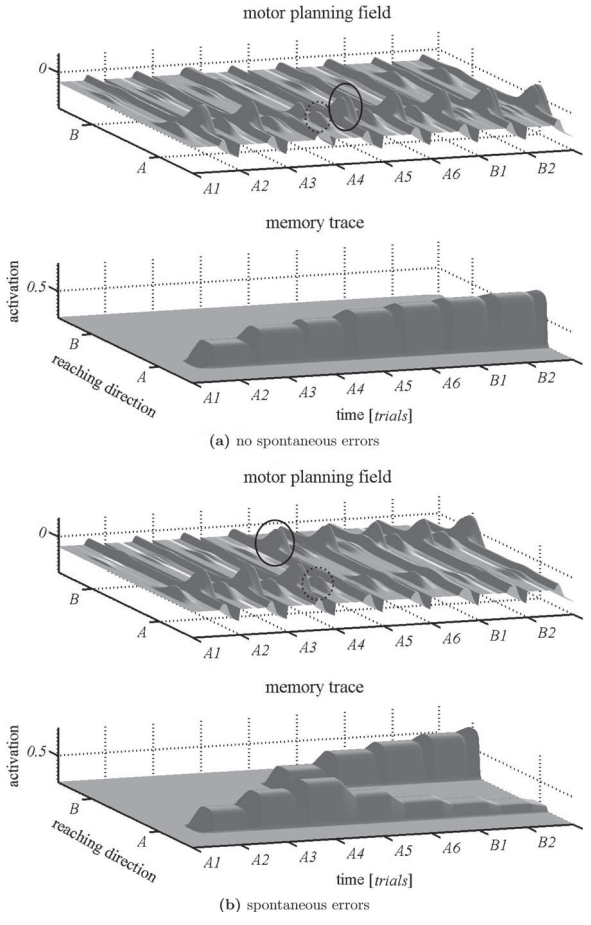

Figure 6. Two simulations of the DFT model in the *A*-not-*B* task at identical parameter settings, differing only in the random fluctuations of the stimulus strength. The movement planning fields (top in (a) and in (b)) and associated memory traces (bottom (a) and in (b)) are shown as functions of reaching direction and time. The positions *A* and *B* of the two target locations are marked along the behavioural dimension, and trials, *A*1 to *B*2, are marked along the temporal dimension. Cue and Reach on the A_4 trials are circumscribed with dashed and, respectively, solid lines. (a) No spontaneous errors and (b) spontaneous errors.

error occurs again on the fifth and sixth *A* trial. This is not surprising, because the memory trace is no longer so strongly biased towards *A*. In fact, the memory trace becomes favourable to *B* by the sixth trial. By the time the cue switches to *B*, there is more input from the memory trace at *B* than at *A*. The model "follows the cue" and "correctly" selects the *B* location on the *B* trials, supported by the memory trace.

This figure thus illustrates how stochastic decisions impact, through the memory trace, on future decisions. We will seek support for this phenomenon by making detailed quantitative analyses of stochastic decisions in a large data set obtained by pooling across a number of studies, and by showing that the model provides a quantitative account for estimated conditional probabilities.

Table 1. Sources from which data is pooled for the analysis of spontaneous errors. Data sizes of the toy version *A*-not-*B* task are listed in the second to last column.

		Datasizes		
Study	Study description	Pool	Toy	$B \neq A$
Smith et al. (1999)	Sit/stand on *A* or *B* trials	62	24	31
Diedrich et al. (2001)	Asymmetric task input	18		
Clearfield et al. (2009)	Specific cue and delay length	108		
Not published	Toy cue and delay length	36	36	
Diedrich, Smith, & Thelen (1998)	Weighted lids	76		34
Data in Table 2	Training trials vs. no-training	17		

Note: When *B* trials differed from *A* trials in procedure, the subset size is listed in the last column.

3. Methods

3.1. Experiments

To identify empirical signatures of the dynamic mechanisms of sensorimotor decision making in infants, we systematically examine the statistics of spontaneous and perseverative errors in experiment and in the model. To this end, we pool data from several published and unpublished studies performed in the Smith and Thelen infant labs at Indiana University, listed in Table 1. From these studies only conditions were retained that (a) tested young 6 to 7 month old infants, (b) applied six *A* and two *B* trials, (c) used a training procedure (Figure 3), and (d) used identical looking lids at the *A* and *B* locations. In total, we analyse the raw data from 317 infants. This includes 50 infants who were tested with hidden toys and 65 infants who were tested with different reaching conditions on *B* than on *A* trials (e.g. standing up vs. sitting). Some experiments included other conditions (such as different types of lids at two locations, Diedrich, Highlands, Spahr, Thelen, & Smith (2001)), but we included only data that conformed with (a)–(d).

In this data set, each infant contributes a total of eight sensorimotor decisions across *A* and *B* trials. This makes it possible to examine the interdependence across these different decisions. In fact, the large number of participants obtained by pooling data across studies made it possible to directly estimate a set of conditional probabilities, that characterise this interdependence among decisions. This novel method is described mathematically as follows.

Probability of spontaneous errors as a function of trial number: On any *A*-trial, *T*, the probability $P[\cdot]$ of a reach to *B* is estimated as the frequency $F[\cdot]$ of reaches to *B* on trial *T* normalised by the population size:

$$P[A_T = B] = \frac{F[A_T = B]}{N}. \tag{1}$$

Persistence of spontaneous errors: On any *A*-trial, *T*, the probability of a second spontaneous error to follow immediately after the first spontaneous error is estimated from the probability to reach to *B* on trial *T*, conditional on the reach on the previous trial $(T-1)$ having been to *B* too and all reaches before that point having been to *A*:

$$\begin{aligned} &P[A_T = B | A_{T-1} = B, A_{T-2} = A, \ldots, A_1 = A] \\ &\quad = \frac{F[A_T = B, A_{T-1} = B, A_{T-2} = A, \ldots, A_1 = A]}{F[A_{T-1} = B, \ldots, A_1 = A]}. \end{aligned} \tag{2}$$

Rate of spontaneous errors: The probability for any infant/simulation to make exactly n spontaneous errors B-on-A is computed from the associated frequency, F, normalised with population size N:

$$P[\#(B\text{-on-}A) \equiv n] = \frac{F[\#(B\text{-on-}A) \equiv n]}{N}. \tag{3}$$

First spontaneous error: The probability that a first spontaneous error occurs on trial number T is estimated as the probability to reach to B on trial T, conditional on all previous reaches having been to A:

$$P[A_T = B | A_{T-1} = A, \ldots, A_1 = A] = \frac{F[A_T = B, A_{T-1} = A, \ldots, A_1 = A]}{F[A_{T-1} = A, \ldots, A_1 = A]}. \tag{4}$$

Perseveration depending on spontaneous errors: The probability of perseverative reaching to A on trial B_1 ($B_1 = A$), conditional on the number, n, of reaches to B on the A trials is estimated from the associated frequencies $F[\,\cdot\,]$:

$$P[B_1 = A | \neq (B\text{-no-}A) \equiv n] = \frac{F[\neq (B\text{-no-}A) \equiv n,\ B_1 = A]}{F[\neq (B\text{-no-}A) \equiv n]}. \tag{5}$$

Comparison of the model to experimental data requires determining values of the model parameters listed in Table A1. We had to handle three kinds of parameters. First, a set of parameters reflects the experimental paradigms that are translated into time courses of the corresponding inputs. The durations of the different phases of each trial, and the numbers of A and of B trials are such parameters. Second, a set of parameters captures essentially descriptive features of the model, which are not critical to the outcomes. For instance, the size of the field, the spatial separation of A and B, the time scale of the activation dynamics and of the memory trace, are parameters that are broadly reflective of what is known about the metrics of reaches and about the time scale of behaviour. Their precise values have little impact on the fit of the measured probabilities. In particular, the fitted probabilities do not reflect response metrics and absolute time. Third, there is a set of parameters that represent the core properties of the neural dynamics and of the role of the different inputs to the field. This includes the parameters that describe the interaction kernel, all input strengths as well as the resting level of the field. This set of parameter values is constrained by postulates about the different dynamic regimes in which the neural dynamics is expected to be in different phases of a trial. For instance, as argued above, the boost that models the pushing of the box into reaching space is expected to push the field through the detection instability, so in that phase of the trial, the field must be in a dynamic regime in which a self-stabilised peak is mono-stable. Restricted by these constraints we chose parameter values in two ways: The parameter values of the interaction kernel were chosen to match those used in the previous model (Thelen et al., 2001, marked by an asterisk in the Table). The parameter values of input strengths and widths were chosen to achieve good experimental fit. This was a fitting process "by hand", in which simulations were performed for different parameter values which were then updated to improve the fit. Note that this fitting did not extensively explore parameter space, but was meant primarily to obtain a rough match to experiment. All widths of input kernels, for instance, were set to 10 (spatial units) for simplicity and input strengths were similarly varied in simple integer ratios (e.g. strength 2 of task input and strength 10 of specific input). One reason to aim at qualitative rather than quantitative fit is that experimental details vary across contributing experiments.

In the simulations, variance derives from the noise sources that are part of the model. These noise source generate fluctuation in time, but also across repetitions of the simulations. To estimate the conditional probabilities for the model, we ran the *A*-not-*B* paradigm a total of 3300 times at fixed values of all parameters except for cue strength that was varied to reflect experimental variations of cuing ($S_{\text{spec}} = 1, 2, \ldots, 11$, with 300 simulations at each value; Clearfield et al., 2009). We thus used a roughly 10 times larger ensemble of repetitions for the model than for experiment to minimise uncertainty in the estimate of the measures taken on the model.

Because non-reaches were not handled uniformly in the different experiments, only individuals who reached on all trials were retained for the analysis. In experiment, this reduced the population from $N = 317$ to $N = 286$. For the model, non-reaches were rare (1 in 300), showing that the model does not capture all reasons for non-reaches.

4. Results

We focus on the quantitative account provided by the model for the stochastic sensorimotor decision making of infants in the *A*-not-*B* paradigm. In each case, we compare the respective probabilities computed from the pooled experimental data with the estimates of these probabilities from model simulations at a fixed set of parameters values. We have also established that the new variant of the DFT model continues to account for all experimental effects that the previous model (Thelen et al., 2001) handled successfully. The quantitative account for the central developmental signature, the age-delay effect (Diamond, 1985), is provided in Appendix 2. Here we focus on how the model accounts for the reaching data of young infants whose stochastic sensorimotor decisions reveal interdependencies across trials.

4.1. Rate of spontaneous errors

Overall, spontaneous errors occur on about 20% of all *A* trials. Figure 7 shows how the rate of spontaneous errors evolves across *A* trials (Equation (1)). For the first four *A* trials, the rate of spontaneous errors actually increases. We will analyse below how this is linked to the training procedure, in which the *A* location is made more attractive for those trials. For trial A_5 and A_6, the rate of spontaneous errors is constant at approximately the overall rate of spontaneous errors. There are two limit cases that predict this constant rate. In one limit case, sensorimotor decisions for *B* on *A* trials are purely stochastic, independent on each trial. In other words, on each trial, there is the same small chance for an infant to reach spontaneously to *B*. In this limit case the probability of making a spontaneous error is independent of previous outcomes of the decision. Figure 7 tests this prediction by including the probability of making a spontaneous error conditional on having made a spontaneous error on the previous trial (Equation (2)). This conditional probability is found to be consistently higher than the overall rate of spontaneous errors, in fact, about twice that rate. This strongly refutes the hypothesis that spontaneous errors are purely stochastic and independent of each other across trials.

In the other limit case, the sensorimotor decision is deterministic across trials. Any given baby systematically either goes to *A* or to *B* on all *A* trials, for instance, due to a side bias.

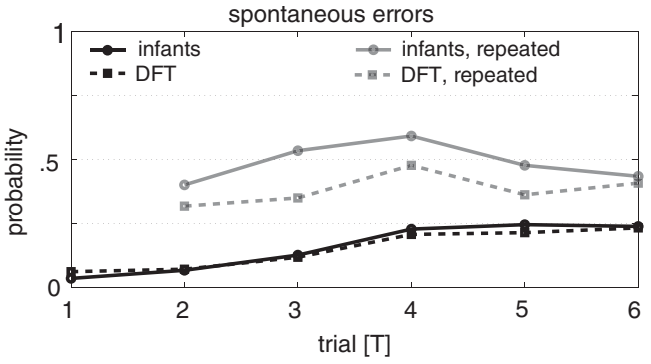

Figure 7. Estimates from experiment (solid lines) and DFT simulations (broken lines) of the rate of spontaneous errors across *A*-trials (black lines). The grey lines show the conditional probability that a reach again goes to *B* on a given *A*-trial given that the first spontaneous reach to *B* has just occurred on the previous trial.

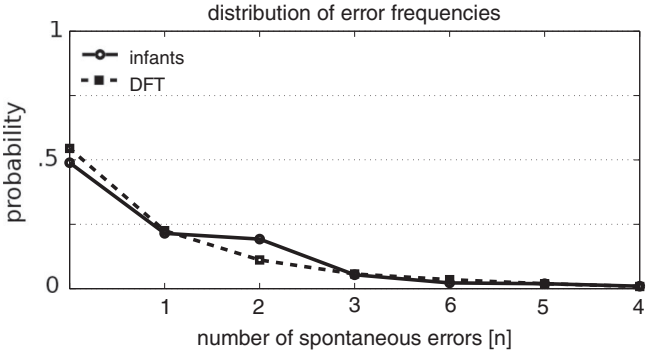

Figure 8. Estimates from infant experiments (solid line) and DFT simulations (broken line) for the probability to make exactly *n* spontaneous errors as a function of *n*.

According to this hypothesis, the overall rate of spontaneous errors reflects the distribution of the side bias across babies and is, therefore, constant across *A* trials. This hypothesis predicts that the conditional probability of repeating a spontaneous error after a previous error should be high (close to one in the limit case of completely deterministic decisions). In fact, this limit case predicts that babies with a bias to *B* should repeat spontaneous errors across the entire *A*-trials phase of the paradigm.

This prediction is tested in Figure 8 showing the probability that an infant/simulation makes exactly *n* spontaneous errors as a function of *n* (Equation (3)). The deterministic account predicts that this probability should have a U-shape: Some infants should systematically make no spontaneous errors, while the biased babies should make a large number of spontaneous errors. Intermediate numbers of spontaneous errors should not be frequent, as these reflect stochastic decision making. The data clearly refute this hypothesis. The monotonic decrease of the probability of *n* spontaneous errors with the number *n* is consistent with a stochastic contribution to sensorimotor decision making.

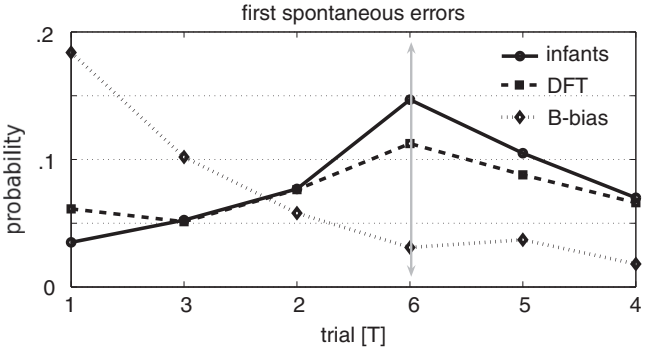

Figure 9. Probability that the first spontaneous error to *B* occurs on trial A_T as a function of trial number *T* for data from infant experiment (solid line), DFT simulation (dashed line), and *B*-bias model (dotted line). The vertical double-arrow marks trial A_4, the first *A*-trial with symmetric task input.

4.2. First spontaneous errors

With these two limit cases out of the way, we recognise that spontaneous errors are interdependent, but also stochastic in nature. How can we understand their evolution in time across *A* trials? To obtain a base line for spontaneous errors, we look at the first occurrence of a spontaneous error across *A* trials (Equation (4)) in Figure 9. This way of analysing spontaneous errors effectively eliminates the influence of previous reaches to *B*. The increase of the probability of a first spontaneous error over the first four *A* trials is reflective of the training procedure and will be examined in detail in a moment. Before we do that, we use this data to refute a more sophisticated version of the side bias account for the data. Such an account is constructed to be consistent with the distribution of infants' spontaneous errors reported in Figure 8. It postulates that all spontaneous reaches come from biased individuals, whose bias strength varies. The strength of the side bias in infants who make *n* spontaneous errors across *A* trials is formalised as a probability, $n/6$ of reaching to *B* on any given *A* trial (6 is the number of *A* trials). A population of infants with different bias strengths is constructed such that they predict the observed rate of making exactly *n* spontaneous errors (in other words, that rate determines the proportion of individuals with bias $n/6$ in the population). Figure 9 includes the rate of first spontaneous errors across *A* trials predicted by this bias model. The predicted probability strongly decreases across *A* trials and clearly is in conflict with the data. We reject, therefore, also this more sophisticated side bias explanation of spontaneous errors.

So now let us look at the remarkable time structure of the probability of the first spontaneous error across *A* trials in light of the training procedure. An intuitive account goes as follows: Initially, spontaneous errors are unlikely, because the asymmetrical presentation of *A* and *B* favour reaches to *A*. Over the first four *A* trials, this bias in the presentation is reduced, making reaches to *B*, that is, spontaneous errors, increasingly likely. After symmetry between *A* and *B* is restored from the fourth *A* trial on, the probability of reaching to *B* decreases, because overall most infants reach to *A* on these trials and thus continue to strengthen the habit of reaching to *A*.

How this intuitive account translates into the DFT model is illustrated in Figure 10 which plots the differences in task input and in input from the memory trace between the *A* and

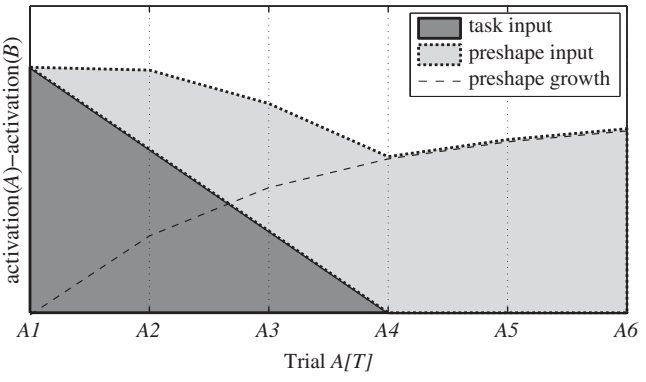

Figure 10. Displayed is the difference of input in the model at the end of the delay between the *A*- and the *B*-location, plotted as a function of *A*-trial number. The accumulated input advantage for *A* (dotted line on top) is the sum of task input advantage for *A* (solid line) and memory trace advantage for *A* (dashed line). These curves were obtained from simulations without noise, in which no spontaneous errors occur.

the *B* location of the field across *A* trials. For the first four *A* trials, the training procedure is modelled by providing more task input at *A* than at *B* (Figure 3). The asymmetry of task input is decreased linearly across the first four *A* trials, reaching zero on A_4. Because we are looking at first spontaneous errors, the memory trace evolves from reaches to *A* only. That memory trace is localised at *A* and grows across consecutive reaches. Together, the two sources of input make for an advantage for the *A* location that decreases over the first four *A* trials and then increases over the remaining two *A* trials. This predicts the pattern of first reaches for the model shown in Figure 9, that is, an increase up to A_4, followed by a decrease.

4.3. A-not-B experiment without training procedure

Clearly, the training procedure impacts on sensorimotor decision making during *A* trials. It is interesting, therefore, to look at variants of the experiment that drop the training procedure. Here we contrast the behaviour of infants exposed to the standard training procedure (Smith et al., 1999) with the behaviour of infants who received six straight *A* trials without any training asymmetry. Both data sets come from a lids-only version of the *A*-not-*B* task. The overall rate of spontaneous errors was 19% with training, but 39% without training.

The high rate of spontaneous errors in the experiment without training procedure makes that performance on *A* trials is close to chance. Looking at the reaches of individual infants reveals, however, that reaches are not randomly distributed across *A* trials (Table 2). Without training, the first reach on A_1 is close to chance level performance, while it is overwhelmingly biased to *A* with training. On subsequent trials, however, infants tend to stick to their initial choice in the no training variant of the experiment. This includes a number of infants who consistently reach to *B* across a large number of *A* trials. With training, spontaneous errors tend to occur in more isolated form in the middle of the *A*-trials phase. This analysis reveals, therefore, the interdependence among subsequent reaching decisions, consistent

Table 2. Data for individual infants in experiment (left column) and for individual DFT simulation runs (right column) are contrasted for the training (top table) versus the no-training (bottom table) condition.

(a) Infants with training

Trail		A_1	A_2	A_3	A_4	A_5	A_6	B_1	B_2
baby	1	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
baby	2	*A*	*A*	*A*	*A*	*A*	*A*	–	*A*
baby	3	*A*	*A*	*A*	*A*	*A*	*A*	–	*A*
baby	4	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
baby	5	*A*	*A*	*A*	*A*	*B*	*A*	*A*	*A*
baby	6	*A*	*A*	*A*	*B*	*A*	*B*	*A*	*A*
baby	7	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
baby	8	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
baby	9	*A*	*A*	*A*	*A*	–	*B*	*A*	*A*
baby	10	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*B*
baby	11	*A*	*A*	*B*	*B*	*A*	*A*	*A*	*B*
baby	12	*A*	*A*	*A*	*A*	*A*	*B*	*A*	*B*
baby	13	*A*	*A*	*A*	*B*	*B*	*A*	*B*	*B*
baby	14	*A*	*A*	*A*	*A*	*A*	*B*	*B*	*B*
baby	15	*A*	*A*	*A*	*A*	*B*	*B*	*B*	*B*
baby	16	*A*	*A*	*A*	*B*	*A*	*B*	*B*	*B*
baby	17	*B*	*A*	*B*	*B*	*B*	*B*	*B*	*B*

(b) DFT with training

Trail		A_1	A_2	A_3	A_4	A_5	A_6	B_1	B_2
siml	1	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	2	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	3	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	4	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	5	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	6	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	7	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	8	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	9	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*B*
siml	10	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*B*
siml	11	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*B*
siml	12	*A*	*A*	*B*	*A*	*A*	*A*	*A*	*A*
siml	13	*A*	*A*	*A*	*B*	*B*	*A*	*A*	*A*
siml	14	*A*	*A*	*B*	*B*	*B*	*B*	*B*	*B*
siml	15	*A*	*B*	*A*	*B*	*B*	*B*	*A*	*B*

(c) Infants, no training

Trail		A_1	A_2	A_3	A_4	A_5	A_6	B_1	B_2
baby	1	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
baby	2	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
baby	3	*A*	*A*	*B*	*A*	*A*	*A*	*A*	–
baby	4	*A*	*A*	*A*	*A*	*A*	*A*	*B*	*B*
baby	5	*A*	*A*	*A*	*A*	*A*	*B*	*B*	*B*
baby	6	*B*	*B*	*A*	*A*	*A*	*A*	*A*	*B*
baby	7	*B*	*B*	*B*	*A*	*B*	*A*	*B*	*B*
baby	8	*B*	*A*	*B*	*B*	*B*	*B*	*B*	*B*
baby	9	*A*	*B*	*B*	*B*	*B*	*B*	*B*	*B*
baby	10	*A*	*B*	*B*	*B*	*B*	*B*	*B*	*B*
baby	11	*B*	*B*	*B*	*B*	*B*	*B*	*B*	*B*
baby	12	*B*	*B*	*B*	*B*	*B*	*B*	*B*	*B*

(d) DFT, no training

Trail		A_1	A_2	A_3	A_4	A_5	A_6	B_1	B_2
siml	1	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	2	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	3	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	4	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	5	*A*	*A*	*A*	*A*	*A*	*A*	*A*	*B*
siml	6	*A*	*A*	*A*	*A*	*A*	*B*	*A*	*A*
siml	7	*B*	*A*	*A*	*A*	*A*	*A*	*A*	*A*
siml	8	*A*	*A*	*B*	*A*	*B*	*A*	*A*	*A*
siml	9	*A*	*A*	*A*	*A*	*B*	*B*	*B*	*B*
siml	10	*A*	*B*	*B*	*A*	*B*	*B*	*B*	*B*
siml	11	*A*	*B*	*B*	*B*	*A*	*A*	*B*	*B*
siml	12	*B*	*B*	*B*	*A*	*B*	*A*	*B*	*A*
siml	13	*B*	*B*	*B*	*B*	*B*	*B*	*B*	*B*
siml	14	*B*	*B*	*B*	*B*	*B*	*B*	*B*	*B*
siml	15	*B*	*B*	*B*	*B*	*B*	*B*	*B*	*B*

Note: Within each sub-table, rows represent individual experimental runs across the trials. The letters "*A*", "*B*", and "–" stand for reaches to *A*, *B*, and non-reaches, respectively. Infants/simulations are ordered from few to many *B* reaches to facilitate the detection of response patterns.

with the role of the memory trace through which sensorimotor decisions affect subsequent sensorimotor decisions.

4.4. Perseverative reaching

Figure 11 shows how the probability to perseveratively reach to *A* on *B* trials (Equation (5)) depends on the number of spontaneous reaches to *B* on *A* trials. Strong perseveration is obtained when no spontaneous errors occur. Because about half of the infants make no spontaneous error (Figure 8), this confirms the prevalence of perseveration. The larger the number of spontaneous errors, the lower the probability of perseverating. In fact, infants with one to four spontaneous reaches to *B* perform at chance level on *B* trials. Infants with 5

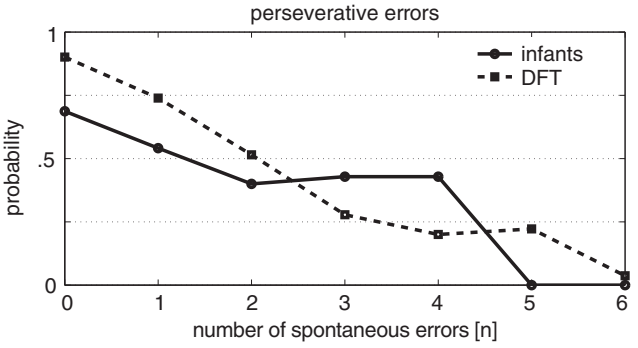

Figure 11. Data from infant experiments (solid line) and DFT simulations (broken line). Shown is the probability of a perseverative error to *A* on the first *B* trial conditional on the number of reaches to *B* on the *A* trials.

or 6 spontaneous reaches to *B* reach correctly on *B*. The model captures a monotonic trend, although the quantitative fit is not perfect.

5. Discussion

We started our examination of sensorimotor decisions from the postulate that such decisions must lead to a stable state that persists long enough in time to steer the generation of motor behaviour while remaining coupled to the sensory surface. Such coupling is observed in online updating of sensorimotor decisions when movement targets shift unexpectedly during movement preparation or generation (Goodale et al., 1986; Prablanc & Martin, 1992; van Sonderen, van der Gon Denier, & Gielen, 1989). Such stable patterns of neural activation are motor plans that differ from alternative motor plans macroscopically. That macroscopic difference persists in time long enough to guide the behaviour. In the DFT model, a self-stabilised peak of activation represents such a motor plan. When the peak is generated in a detection instability, it amplifies a potentially microscopic difference between alternative choices into a macroscopic difference.

The *A*-not-*B* paradigm entails a sensorimotor decision in the form of a delayed response task. The "go" signal is separated in time from the cue. We focussed on the fact that young infants do not always follow the cue in the *A*-not-*B* paradigm, not only on the *B* trials, when they perseverate, but also on *A* trials, when they make spontaneous errors, reaching to *B* when *A* was cued. The decisions are thus partially stochastic, and this brings the issue of whether the decisions are based on macroscopic rather than microscopic differences into the foreground.

The classical literature on movement preparation (reviewed in Erlhagen & Schöner, 2002) and much of the recent neurophysiological work on sensorimotor decisions (Cisek & Kalaska, 2005) made used of reaction time paradigms, in which movements are initiated as soon as an imperative signal is processed. Those paradigms invite accounts for decision making that focus on the gathering evidence for one choice over other possible choices. The classical diffusion model (reviewed in Ratcliff et al., 1999), for instance, postulates that evidence is accumulated along a decision dimension in time until a criterion level is reached.

In a sense, this is a mechanism to generate a macroscopic difference between the neural representation of the different choices (Gold & Shadlen, 2007). The decision itself is triggered by an event that ends this process of integration. The diffusion model does not address in which way the neural representation of the decision persists and thus does not lend itself to an account for delayed response tasks.

The DFT account for movement preparation (Cisek, 2006; Erlhagen & Schöner, 2002) provides an alternative. In that theoretical framework, sensorimotor decisions are represented by self-stabilised peaks of activation that arise within fields defined over relevant movement parameters. In reaction time paradigms, the imperative signal provides localised input along those dimensions and pushes the activation field through the detection instability, at which a self-stabilised peak forms. Other inputs such as precues to upcoming movement targets, perceptual inputs that reflect the task set, or input from a memory trace of earlier sensorimotor decisions may contribute to the formation of the peak. The influence of the memory trace is typically minor, primarily affecting reaction time, not the actual motor response. The uncertainty about the sensorimotor decision can be enhanced by "forcing" the decision early in the selection process. This can be done in the timed movement initiation paradigm (Ghez et al., 1997), in which a metronome provides the go signal at a stimulus-response time interval that can be varied experimentally. For very short stimulus-response intervals, the role of inputs other than the imperative signal is amplified, leading to "default" responses to the most frequent or the average movement direction, for instance.

Infants reveal such influences in the delayed response paradigm in the sense that the confluence of various sources of inputs impacts on their sensorimotor decision. Our theoretical analysis helps to understand, why this is the case. Young infants decide "late", that is, they only form a motor plan at the end of the delay, when the reaching locations come into range. Young infants do not form a stable motor plan when first cued to a reaching location and cannot keep such a plan in working memory. The cue is therefore a weak contribution to their sensory motor decision, rather than an imperative signal, and this is why they do not necessarily follow the cue. The cue plays, therefore, a relatively minor role in specifying the outcome of the sensorimotor decision, as attested by the model's account for the training procedure, which may largely be responsible for what looks like "following the cue" on *A* trials, at least in the lids-only variant of the *A*-not-*B* task that we focussed on. In fact, without the training procedure, infants respond close to chance level on *A* trials. Thus, the account we provided in this paper shows how reaching decisions of young infants in the *A*-not-*B* task may provide a window into sensorimotor decision making that complements the more typical reaction time paradigms used in adult and animal research.

The earlier DFT model (Thelen et al., 2001) did not capture this fact, that young infants decide late in the trial. In that earlier model, a peak is formed in response to the cue, but decays during the delay. It is that earlier peak that determines which memory trace is laid down. That early decision is dominated by the cue. In the earlier model, therefore, the memory trace reflects the cue rather than the final sensorimotor decision. That decision was based on a read-out procedure in which the difference in levels of activation between the two choices was typically microscopic, and thus did not leave a memory trace. The same is true about the PDP model of perseveration (Munakata, 1998), which does not provide an actual neural mechanism for the sensorimotor decision itself.

More generally, the "naive" form of "winner takes all" at read-out may never be a good account for sensorimotor decisions, because it does not guarantee that a macroscopic

state of neural activation is stabilised against distractors and persists sufficiently in time to generate motor behaviour and to leave a memory trace. Note that this fact is less obvious in reaction time paradigms only because they typically profit from a strong imperative stimulus that disambiguates responses even on a short time scale. When fast responses are given without out a clear imperative stimulus (such as saccadic eye movements or naturalistic action sequences), however, the underlying stability problem becomes obvious, a point to which we will return in a moment.

We have provided evidence for the interdependence between subsequent sensorimotor decisions in infant perseverative reaching through a series of conditional probabilities that probed the stochastic nature of reaching behaviour. We were able to exclude both limit cases, purely stochastic decision making that would be independent across repeated reaching on subsequent trials, but also deterministic reaching within each individual, biased either to *A* or to *B*. Thus, spontaneous errors, reaches to *B* following a cue at *A*, were shown to predict a higher than average probability to repeat that decision on subsequent *A* trials and a lower than average probability to reach to *A* on *B* trials. This form of interdependence arises, in the DFT account, from the memory trace that each reaching decision lays down.

Other developmental evidence for the memory trace comes from a metric attraction to previous hiding locations in the sandbox version of the *A* not *B* task (Schutte & Spencer, 2002; Schutte, Spencer, & Schöner, 2003), a related form of attraction of spatial working memory items to spatial locations that were previously memorised (Schutte & Spencer, 2009), as well as perseveration in task switching contexts such as the Wisconsin card sorting task (Buss & Spencer, 2014; Morton & Munakata, 2002). The memory trace was originally postulated as an account for a large number of set effects in reaction time tasks (Erlhagen & Schöner, 2002). For instance, the dynamics of the memory trace accounted for how the probability of choices affects response time in Hyman's law. In the timed movement initiation paradigm, the preshaping of the neural representation of movement plans that reflects previous motor decisions can be directly observed in the form of the "default" distribution of motor responses generated at very short stimulus-response intervals (Erlhagen & Schöner, 2002; Ghez et al., 1997).

Still, intertrial effects are not very prominent in the adult movement preparation literature. One reason is again that most of these studies have a strong imperative stimulus, that disambiguates the sensorimotor decision while also driving response initiation. The influence of the memory trace is comparatively weak and may not be observable in many cases. Note that the parameter values we found useful to account for the influence of prior reaches was five times smaller than the strength of the cue (2 vs. 10), consistent with this explanation. If the imperative signal is removed or weakened, intertrial effects are, in fact observed. This is observed for saccadic eye movements, for instance (Fecteau & Munoz, 2003; Genovesio & Ferraina, 2014). Such competitive selection decisions also highlight the need to actively inhibit distractor movement targets, and thus underline the need for stability. Only if the neural representation of target selection is stable may that representation be updated online to target shifts without switching to a distractor target. Similar observations may be relevant to action slips in naturalistic sequences of actions (Botvinick & Bylsma, 2005).

How do these insights into sensorimotor decisions and their interdependence impact on the developmental account for perseverative reaching? In the original DFT model, development was postulated to be reflected in a shift from a primarily input-driven dynamic regime for younger infants to a primarily interaction-dominated regime for older infants (Thelen

et al., 2001). In the input dominated regime, activation patterns are uniquely determined by input, while in the interaction dominated regime, activation patterns are not necessarily uniquely determined by input. This notion has since been generalised to a wide set of developmental changes under the label of the "spatial precision hypothesis" (Schutte et al., 2003; Simmering, Schutte, & Spencer, 2008) reflecting a concomitant narrowing of the spatial range of the stabilising neural interaction. In the variant of the DFT model we presented here, even young infants will have a self-stabilised peak of activation when they initiate a reach. So they are capable of engaging sufficient neural interaction to bring this about. What does this mean then for the developmental account? There is no fixed relationship between neural connectivity and dynamic regime. Different stages of development are thus not characterised by particular dynamic regimes. Instead, as the neural connectivity that shapes the neural dynamics develops, that change affects the range of conditions under which any particular dynamic regime can be reached. Specifically, young infants do not generate self-stabilised motor plans unless there is sufficient convergent stimulation. In the *A*-not-*B* paradigm this is only the case when the reaching targets are within the spatial range of infants reaches. This is why the younger infants form a reaching decision only at the end of the delay. Older infants, in contrast, are capable to stabilising movement plans with less convergent input, so that they form a movement plan following the cue, and in many cases sustain that movement plan through the delay. This is consistent with the account for perseveration in toddlers in the sandbox version of *A*-not-*B* (Schutte et al., 2003). In that account, toddlers generate a self-stabilised peak when cued, that then drifts towards locations pre-activated by the memory trace during the delay in the absence of any visual markers that would lock the peak in place.

In fact, the memory trace itself may push infants into the dynamic regime in which they may make stable reaching plans earlier in the trial! This idea has been proposed and tested to account for the emergence of visual working memory across a number of repeated presentations of the same stimulus (Perone & Spencer, 2013). Although the memory trace appears to reduce infants' flexibility when the target of an action changes, it is also a mechanism how sensorimotor decisions become less dependent on supportive sensory information and more endogenously driven (Schöner & Dineva, 2007; Thelen & Smith, 1994).

6. Conclusions

Different outcomes of a sensorimotor decision are associated with macroscopic differences in neural activation pattern. These differences lead to different learning trajectories so that sensorimotor decision have consequences beyond the immediate behavioural act. Theoretical accounts that relegate decisions to a process of "reading out" neural activation patterns miss this aspect of decision making. Because there is no reason why this fundamental issue should be specific to sensor-motor decisions, we postulate that all neural accounts of decision making must integrate the decision process into the underlying neural dynamics and take into account the implications for memory formation and learning.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was started in collaboration with the late Esther Thelen, whom we fondly remember and in whose debt we are. Initial work on this research was funded by the National Institutes of Health grant [R01 HD 22830] to Esther Thelen, the International Graduate School of Neuroscience of Ruhr-Universiät-Bochum, Germany [IGSN], and EU Commission [NeuralDynamics]. We thank John P Spencer and Linda Smith for many fruitful discussions, and multiple critical readings of the manuscript.

ORCID

Gregor Schöner http://orcid.org/0000-0002-4776-9998

References

Amari, S. I. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*, 77–87.

Bicho, E., Mallet, P., & Schöner, G. (2000). Target representation on an autonomous vehicle with low-level sensors. *The International Journal of Robotics Research*, *19*, 424–447.

Botvinick, M., & Bylsma, L. (2005). Distraction and action slips in an everyday task: Evidence for a dynamic representation of task context. *Psychonomic Bulletin & Review*, *12*(6), 1011–1017.

Buss, A. T., & Spencer, J. P. (2014). The emergent executive: A dynamic field theory of the development of executive function. *Monographs of the Society for Research in Child Development*, *79*(2), 1–103.

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: A computational model. *Journal of Neuroscience*, *26*(38), 9761–9770.

Cisek, P., & Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. *Neuron*, *3*(45), 801–814.

Clearfield, M. W., Dineva, E., Smith, L. B., Diedrich, F. J., & Thelen, E. (2009). Cue salience and infant perseverative reaching: Tests of the dynamic field theory. *Developmental Science*, *12*(1), 26–40.

Diamond, A. (1985). Development of the ability to use recall to guide action, as indicated by infants' performance on A-not-B. *Child Development*, *56*, 868–883.

Diedrich, F. J., Highlands, T. M., Spahr, K. A., Thelen, E., & Smith, L. B. (2001). The role of target distinctiveness in infant perseverative reaching. *Journal of Experimental Child Psychology*, *78*, 263–290.

Diedrich, F. J., Smith, L. B., & Thelen, E. (1998). *Motor memory influences performance in the A-not-B task*. Paper presented at the international conference on infant studies, Atlanta, GA.

Erlhagen, W., Bastian, A., Jancke, D., Riehle, A., & Schöner, G. (1999). The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. *Journal of Neuroscience Methods*, *94*(1), 53–66.

Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*, 545–572.

Fecteau, J. H., & Munoz, D. P. (2003). Exploring the consequences of the previous trial. *Nature Reviews Neuroscience*, *4*, 435–443.

Genovesio, A., & Ferraina, S. (2014). The influence of recent decisions on future goal selection. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1655), 20130477.

Georgopoulos, A. P. (1990). Neurophysiology of reaching. In M. Jeannerod (Ed.), *Attention and performance* (pp. 227–263). Hillsdale, NJ: L. Erlbaum Associates.

Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989). Mental rotation of the neural population vector. *Science*, *243*, 1627–1630.

Ghez, C., Favilla, M., Ghilardi, M. F., Gordon, J., Bermejo, R., & Pullman, S. (1997). Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental Brain Research*, *115*, 217–233.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*(1), 535–574.
Goodale, M. A., Pélisson, D., & Prablanc, C. (1986). Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*, *320*, 748–750.
Kloeden, P. E., & Platen, E. (1999). *The numerical solution of stochastic differential equations* (2nd ed.). Berlin Heidelberg: Springer-Verlag.
Mareschal, D. (1998). To reach or not to reach ... that is the question. *Developmental Science*, *1*(2), 198–199.
Morton, J. B., & Munakata, Y. (2002). Active versus latent representations: A neural network model of perseveration, dissociation, and decalage. *Developmental Psychobiology*, *40*(3), 255–265.
Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: A PDP model of the AB task. *Developmental Science*, *1*(2), 161–184.
Munakata, Y., McClelland, J. L., Johnson, M. J., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, *104*(4), 686–713.
Perone, S., & Spencer, J. P. (2013). Autonomy in action: linking the act of looking to memory formation in infancy via dynamic neural fields. *Cognitive Science*, *37*(1), 1–60.
Piaget, J. (1954). *The construction of reality in the child*. New York, NY: Basic Books.
Prablanc, C., & Martin, O. (1992). Autonomous control during hand reaching at undetected two-dimensional target displacements. *Journal of Neurophysiology*, *67*, 455–469.
Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261–300.
Rumelhart, D. E., McClelland, J. L., & The PDP Research Group (1986). *Parallel distributed processing–volume 1: Foundations*. Cambridge, MA: The MIT Press.
Schöner, G., & Dineva, E. (2007). Dynamic instabilities as mechanisms for emergence. *Developmental Science*, *10*(1), 69–74.
Schöner, G., Spencer, J. P., & the DFT research group (2016). *Dynamic thinking: A primer on dynamic field theory*. New York, NY: Oxford University Press.
Schutte, A. R., & Spencer, J. P. (2002). Generalizing the Dynamic Field Theory of the A-not-B error beyond infancy: Three-year-olds' delay- and experience-dependent location memory biases. *Child Development*, *73*(2), 377–404.
Schutte, A. R., & Spencer, J. P. (2009). Tests of the dynamic field theory and the spatial precision hypothesis: Capturing a qualitative developmental transition in spatial working memory. *Journal of Experimental Psychology. Human Perception and Performance*, *35*(6), 1698–1725.
Schutte, A. R., Spencer, J. P., & Schöner, G. (2003). Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, *74*(5), 1393–1417.
Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. *Brain Research*, *1202*, 68–86.
Smith, L. B., Thelen, E., Titzer, R., & Mclin, D. (1999). Knowing in the context of acting: The task dynamics of the A-not-B error. *Psychological Review*, *106*(2), 235–260.
van Sonderen, J. F., van der Gon Denier, J. J., & Gielen, C. C. A. M. (1989). Motor programmes for goal-directed movements are continuously adjusted according to changes in target location. *Experimental Brain Research*, *78*, 139–146.
Spencer, J. P., & Perone, S. (2008). Defending qualitative change: The view from dynamical systems theory. *Child Development*, *79*(6), 1639–1647.
Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences*, *24*(1), 1–34.
Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: The MIT Press.
Wellman, H. M., Cross, D., & Bartsch, K. (1987). Infant search and object permanence: A meta-analysis of the A-not-B error. *Monographs of the Society for Research in Child Development*, *54*(3), 1–67.

Appendix 1. DFT: mathematical formulation

Here we list the model equations used for the DFT simulations. The parameter values are specified in Table A1.

Field dynamics: The field activation evolves as described by the dynamical system:

$$\tau_{\text{mp}}\dot{u}_{\text{mp}}(\psi,t) = -u_{\text{mp}}(\psi,t) + h_{\text{rest}} + \cdots \text{inputs} + \text{interactions} + \text{boost} + \text{noise}. \tag{A1}$$

Here, $\dot{u}_{\text{mp}}$ is the rate of change of the field activation, which is inversely proportional to the current level of activation, $u_{\text{mp}}(\psi,t)$. This generates stability: At large activation levels the rate of change becomes negative, pulling activation back down. At large negative activation levels, the rate of change is positive, pulling activation back up. In the absence of other influences, the resting level, $h_{\text{rest}} < 0$, is the stable, spatially and temporally constant solution, $u_{\text{mp}} = h_{\text{rest}}$. The parameter τ_{mp} determines the time scale of evolution.

Environmental inputs: Perceptual structure in the direction ψ_{loc} or a cue towards movement direction ψ_{loc} are modelled as Gaussian functions of the form

$$\mathcal{G}_{\text{loc}}(\psi,t) := S_{\text{loc}} \cdot \exp\left[-\frac{(\psi-\psi_{\text{loc}})^2}{2\cdot\alpha_{\text{loc}}^2}\right], \tag{A2}$$

where S_{loc} is a measure of the strength of the input and α_{loc} determines its angular range. To reflect the time structure of stimulation, the input strength is modelled as a time-varying function with

Table A1. Parameters used in the DFT simulation.

Parameter name	Value [units]	Description
`delta`	1 [10 ms]	Euler step Δ
`tau`	100 [10 ms]	τ_{mp}, time scale of u_{mp}
`relax`	300 [10 ms]	Prerelaxation phase duration
`specific`	300 [10 ms]	Cueing phase duration
`delay`	300 [10 ms]	Delay phase duration
`off`	700 [10 ms]	Reaching phase duration
`trials_a`	6 [number]	*A* trials
`trials_b`	2 [number]	*B* trials
`fieldsize`	1–201	Field units ψ, discretisation
`loc_a`	50 [field unit]	*A* position
`loc_b`	150 [field unit]	*B* position
`h_rest`	−12 [activation]	Resting level h_{rest}
`h_boost`	9 [activation]	Boost h_{boost}
`sigma_w`*	10 [field units]	Excitatory range width σ
`w_exite`*	2 [activation]	Excitatory strength ω_e
`w_inhib`*	−1 [activation]	Inhibitory strength ω_i
`thrh_u`*	0 [activation]	Threshold f (the turning point)
`beta`*	1.5	β, slope of the sigmoidal threshold f
`q_task`	0.1 [activation]	Noise strength
`task_str`	2 [activation]	Task input strength S_{task}
`task_all(trial)`	[1.2, 0.8, 0.4, 0, . . . , 0]	$F_{A/B}$(trial), asymmetry advantage for *A*
`task_width`	10 [field units]	Task input range (Gaussians $\mathcal{G}$ width)
`spec_str`	10 [activation]	Specific input strength S_{spec}
`spec_width`	10 [field units]	Specific input range ($\mathcal{G}$s width)
`pre_str`	1 [activation]	Preshape strength S_{pre}
`tau_pre`	1000 [10 ms]	τ_{pre}, time scale of u_{pre} ($\tau_{\text{pre}} = 4\cdot\tau_{\text{mp}}$)
`pre_beta`	1.5	Slope of the preshape threshold f_{pre}

Note: Parameters marked with * are used for the interaction and threshold functions which are defined as in the original model (Thelen et al., 2001).

on and off periods. Input is also obtained from the memory trace, which we will discuss below (Equation (A7)).

Neural interactions: Mathematically, interactions are contributions to the rate of change of activation at a given field site, ψ, from all other sites, ψ':

$$g_{\text{coop}}(\psi, t) := \int \omega_{\text{coop}}(\psi - \psi') \cdot \sigma[u_{\text{mp}}(\psi', t)]\, \mathrm{d}\psi'. \tag{A3}$$

The integral sums over the (continuously many) contributions of all sites, ψ'. Only sites, ψ', with sufficient levels of activation contribute. This is described mathematically by applying a non-linear threshold function, $\sigma(u)$, to the activation, a so-called "sigmoid':

$$\sigma[u_{\text{mp}}(\psi', t)] := \frac{1}{1 + \exp(-\beta_{\text{coop}} \cdot u_{\text{mp}}(\psi', t))}. \tag{A4}$$

This threshold function takes values between zero and one. The "synaptic" strength, $\omega_{\text{coop}}(\psi - \psi')$, of each contribution depends on the distance between ψ and ψ':

$$\omega_{\text{coop}}(\psi - \psi') := -\omega_{\text{inhibit}} + \omega_{\text{excite}} \cdot \exp\left(-\frac{(\psi - \psi')^2}{2 \cdot \alpha_{\text{coop}}^2}\right). \tag{A5}$$

Fluctuating input: Noise in the field comes from letting the input strengths S_{loc} in Equation (A2) fluctuate:

$$\tilde{S}_{\text{loc}} = S_{\text{loc}} + q_{\text{noise}} \cdot \eta_{\text{loc}}(t). \tag{A6}$$

where $\eta_{\text{loc}}(t) \sim N(0, 1)$ is a normally distributed random variable. Its contribution to input strength is scaled with a constant factor $q_{\text{noise}} = 0.1$.

Memory trace: Mathematically, the memory trace is generated as a simple low-pass filter dynamics that takes place in a second activation field defined over the same behavioural dimension as the motor planning field, from which it receives input:

$$\tau_{\text{pre}} \cdot \dot{u}_{\text{pre}}(\psi, t) = \big(- u_{\text{pre}}(\psi, t) + \sigma[u_{\text{mp}}(\psi, t)]\big) \cdot \mathcal{X}_{\text{act}}(t). \tag{A7}$$

Only sufficiently activated locations of the motor planning field contribute. Here the input is normalised to the interval from 0 to 1 by the sigmoidal threshold function σ (Equation (A4)). The memory trace tracks self-stabilised (suprathreshold) peaks in the motor planning field on a slower time scale, $\tau_{\text{pre}} \ll \tau_{\text{mp}}$. The function $\mathcal{X}_{\text{act}}(t)$ indicates that the memory trace is updated only when a reach is possible, here approximated by the time window during which the field is boosted into the response mode. While reaching is not possible because the box is outside reaching space, the memory trace remains unchanged.

Task input: The presence of the two lids – the task input – is formalised as the sum of two Gaussians

$$s_{\text{task}}(\psi, t) = [\tilde{S}_{\text{task}} + F_{A/B}(t)] \cdot \mathcal{G}_A(\psi) + \tilde{S}_{\text{task}} \cdot \mathcal{G}_B(\psi) \tag{A8}$$

with mean strength $S_{\text{task}} = 2$ that fluctuates (Equation (A6)). Stronger task input at the *A* location on the first three trials is modelled by adding an asymmetric term $F_{A/B} = 1.2$, 0.8, and 0.4, respectively.

Specific cue: During the cuing phase, the task input is replaced by a specific input at the cued site

$$s_{\text{spec}}(\psi, t) = \begin{cases} \tilde{S}_{\text{spec}} \cdot \mathcal{G}_A(\psi) & \text{on } A \text{ trials,} \\ \tilde{S}_{\text{spec}} \cdot \mathcal{G}_B(\psi) & \text{on } B \text{ trials,} \end{cases} \tag{A9}$$

which has an average amplitude of $S_{\text{spec}} = 10$ that fluctuates (Equation (A6)). The strong specific cue induces a relatively strong peak in the motor planning field which decays, however, during the delay so that the field relaxes to the activation levels induced by the temporally persistent task and memory trace inputs.

Homogeneous boost input: At the end of the delay, a homogeneous boost $h_{\text{boost}} = 9$ of activation is added at all field sites. The boost, together with the other inputs and fluctuations, typically lifts

the field activation sufficiently close to the cooperativity-threshold that a self-stabilised peak forms. Its position represents the reaching location. In the context of the *A*-not-*B* task, this is either the *A* or the *B* location, the only sites receiving localised input. Non-reaches may occasionally occur when activation remains below the threshold even after boosting, so that the inputs are not sufficient to stabilise a reaching decision.

Numerics: Simulations of the model are realised using standard numerical procedures for stochastic differential equations, in particular, the Euler iteration (Kloeden & Platen, 1999). Probabilistic results are generated from multiple simulation runs.

Experimental paradigms are modelled by adjusting the relative strengths, positions, or durations of input presentations.

Appendix 2. Replicating the classical age-delay interaction

Here we test the extended DFT model on Diamond's classical finding of an age-delay interaction (Diamond, 1985). This confirms that the new model preserves quantitative results demonstrated by Thelen et al. (2001), using the same developmental hypothesis, namely that neural interactions strengthen over development.

A.1. Perseveration depends on the delay

In the model, the main difference between *A* and *B* trials is the extent to which inputs converge. On *A* trials, the task input, the specific cue, and the memory trace all favour the *A* location. On *B* trials, the specific cue to *B* competes with the memory trace favouring *A*. Consequently, perseveration should decrease if the memory trace is weakened or if the specific cue is strengthened. The former will be addressed below by looking at the role of behavioural history. The latter is a consequence of the time continuous dynamics, Equation (A1). Activation induced by the specific cue decays gradually after cue removal. The cue's impact can thus be strengthened by increasing its attractiveness or by reducing the delay so that its activation does not decay as strongly (Clearfield et al., 2009). This dynamical property of the motor planning field accounts for the classical finding of a gradual reduction of perseveration with a gradual decrease of the delay (Diamond, 1985).

To understand how the cue-delay interaction works we need to show how a graded reduction of the level of activation at the cue-induced peak at *B* translates into an increased probability of making a perseverative reach, that is, of generating a self-stabilised peak at the alternate *A* location. The key element is that the non-linear neural interaction within the dynamic field may translate fluctuating input (Equation (A6)) into a stable selection decision. This decision occurs as the field goes through an instability, for example, when the homogeneous boost drives the field from an input-driven to an interaction-driven regime.

How this key mechanism leads to a graded dependence of perseveration on the delay is illustrated in Figure A1, which shows the average activation pattern for different delay durations. Higher advantages of subthreshold activation translate into a higher selection probability. Given that shorter delays lead to more activation at the cued side, this explains how the probability of selecting the cued *B* location increases with shorter delays.

The strength of the cue determines the level from which the peak decays. For large cue strength, responses are correct despite a delay. For weak cue strengths, perseveration occurs even in the absence of a delay. Both limit cases have been observed experimentally as well as quantitatively accounted for by the DFT model (Clearfield et al., 2009).

A.2. Developmental trajectory

Thelen et al. (2001) proposed that in older infants, the neuronal dynamics that generates movement plans is capable of sustaining an activation peak induced by the specific cue through the delay period. This change over development was hypothesised to reflect a gradual increase in the relative strength of cooperative neuronal interaction and was modelled simply by assuming that the resting level in

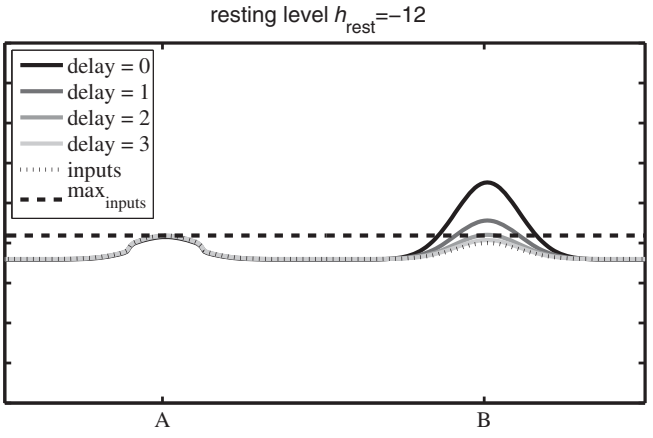

Figure A1. Competition of specific cue and preshape activation in the motor planning field during the delay period of the first *B* trial in a simulation of the young model ($h_{\text{rest}} = -12$) without noise. Snapshots of the motor planning field at different fixed times (grey-scale coded, cf. legend) demonstrate the decay of cuing activation. Without delay, the strong peak at *B* (darkest solid curve) is induced by the just-recent cue presentation. This activation decays (lighter solid curves) and converges to the task and preshape inputs (dotted curve) that are persistent during the delay. Because of previous reaches to *A*, the maximum of preshape plus task input is at *A* (it is marked by the dashed line for better comparison with the decaying peak at *B*). The field activation at the recently cued *B* location fades away during the delay and eventually drops slightly below the preshape plus task input maxim at *A* (lightest solid curve at *B* is below the dashed line). How the differences of activation at *A* and *B* translate into probabilities to reach at either location is explained in the text.

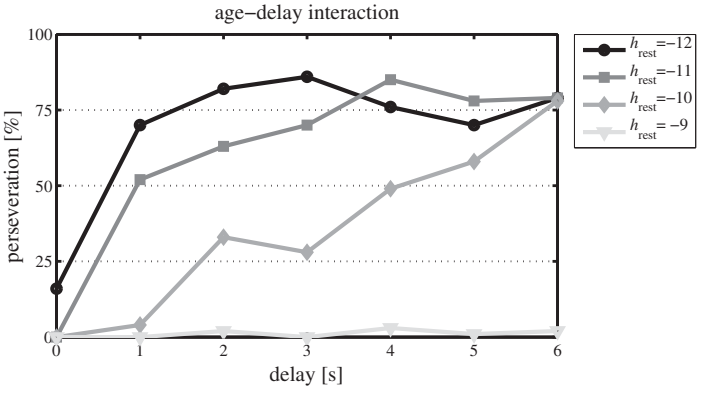

Figure A2. Age-delay interaction in the *A*-not-*B* experiment. The probability of perseveration (*y*- axis) is plotted against different delays (*x*-axis) and resting levels (grey-scale coded: larger values model older infants).

the field increased with age. Increased resting level makes it easier for activation to pass the threshold a peak. This is also how we account for development in the new model.

To examine how the probability of a perseverative error depends on age and delay, increasing resting levels are systematically tested against increasing delays, 100 simulation runs for each combination. From this, the probability of perseveration is computed as the percentage of reaches to the *A* location on the first *B* trial. Figure A2 shows that the model accounts for the empirically established

age-delay interaction in which older infants begin to perseverate at longer delays (Diamond, 1985; Wellman et al., 1987). Reaching is correct at all ages/resting levels in the absence of a delay. For increasing delay, the probability of perseveration increases more strongly for the "younger" models (at low resting levels, $h_{\text{rest}} = -12$ or -11) than for "middle aged" models ($h_{\text{rest}} = -10$). In all cases, the peak induced by the specific cue at *B* decays to some extent, leading to increased probability of reaching to *A* instead. This decay is slower, the stronger self-stabilising interaction is engaged. The "older" model ($h_{\text{rest}} = -9$) does not perseverate at any delay because the peak induced by the specific cue is stably sustained during the delay.

The critical difference between the "younger" and "older" infant models lies in the conditions under which either may first form and sustain self-stabilised peaks of activation. Forming a peak is difficult for the "young" model, which remains near the threshold at which interaction engages both when specific input is provided as well as when the boost is applied. While it passes through this critical point, the system is sensitive to fluctuations and changes in input. This explains why reaching decisions made by the "young" infant model depend strongly on variations of the task context. This also predicts that younger infants show more individual differences even when performing in the same task context, as small variations in their neuronal dynamics may shift them above or below threshold. In the "older" infant model, in contrast, the higher resting level makes it easier for activation to pass and remain beyond the threshold. The system spends less time near threshold and is less sensitive to perturbations and variations of task context. In summary, the developmental trajectory can be accounted for by a graded change of the neuronal dynamics from which a transition from an input-dominated to an interaction-dominated regime emerges (see also Spencer & Perone, 2008).