

Reaching for objects: a neural process account in a developmental perspective

Gregor Schöner and Jan Tekülve and Stephan Zibner
Institut für Neuroinformatik
Ruhr-Universität Bochum
Germany

September 18, 2017

1 Introduction

Reaching and grasping objects is an evolutionarily late achievement. Throughout the animal kingdom, much object oriented action is achieved with the whole body, in mammals often with the head and snout. Because the vision sensor is anchored in the head, such object-oriented movements can be achieved with simple control strategies, such as visual servoing (Ruf & Horaud, 1999), that make limited demands on perception, estimation, and movement planning. Reaching with an actuator that is separate from the main visual sensor is prevalent in primates. Humans excel at object oriented manipulation tasks, much exceeding the skills of other primates. This is a developmental achievement as witnessed by the long and intense period of learning to reach (von Hofsten, 1984, 1991; Thelen, Corbetta, & Spencer, 1996; Berthier & Keen, 2006)

Spatial orientation and navigation may be achieved based on unsegmented visual information (Schöne, 1984; Hermer & Spelke, 1994). In contrast, reaching makes considerable demands on object perception. (1) To reach and grasp an object, that object must be visually segmented against the background and its pose be estimated. (2) Spatial information about the object must be transformed from a visual reference frame into a reference frame in which motor commands to the hands may be formed. Such coordinate transforms are computationally demanding (Pouget & Snyder, 2000; Schneegans, 2015) and learning them is a challenge (Chao, Lee, & Lee, 2010; Sandamirskaya & Conradt, 2013). (3) Motor commands must be generated that drive the hand toward the object and bring it into contact with the object with a small enough terminal velocity that enables grasping. This is particularly challenging as infants are weak relative to the mass of their limbs. Because the force/weight relationship changes during growth, the motor commands must be updated over development. (4) Generating and controlling successful reaches also entails solving the degree of freedom problem, that is, distributing to the many muscles that contribute to arm movement a

coordinated set of commands that achieve contact with the hand (Sporns & Edelman, 1993; Latash, 2008). Little is known to date about how this problem is solved during development. (5) Finally, reaching movements must be initiated and then terminated when the hand makes contact with the object. Typically, successful reaches and grasps are followed by further actions like mouthing or banging the object on a table. The sequential organization of such component movements is part of the motor skill involved in reaching.

This broad range of component processes that span perception, cognition, and motor control, may be a reason why many questions about object-oriented reaching and grasping behavior remain unanswered (see e.g., section on motor control in Lisman (2015)). At one level, a developmental perspective may be particularly attractive. In order for an infant to successfully reach and grasp an object, all processes must be in place and coordinated with each other. Observing how reaching is assembled over development may help understanding how the pieces come together.

Alas, studying the development of reaching empirically is very hard. Babies do not follow task instructions. They must be seduced to repeatedly reach and parametric manipulations of the reaching movements are difficult to impose. Reaches are highly variable and individual differences important (Thelen et al., 1996). It is hard to measure movement trajectories in detail, as preparing infants with markers takes time and patients, inducing attrition. Once movement trajectories have been obtained, their analysis is hampered by the difficulty to reliably detect onsets and offsets of individual movements and thus, to align trajectories across trials (see, however, (Corbetta & Thelen, 1995) for a systematic trial-by-trial solution). In fact, young infants often move continuously, sometimes in approximately periodic fashion.

To date, three data sets about the development of reaching are most elaborate (von Hofsten, 1991; Thelen et al., 1996; Berthier & Keen, 2006). Their properties are reviewed in detail and commented upon in (Caligiore, Parisi, & Baldassarre, 2014). Unfortunately, the data are not strongly constraining for theoretical accounts as of now, even though they include longitudinal studies that provide samples of reaching behavior in individuals along developmental time. A principle problem of all accounts of motor learning during development is, of course, that the actual process of motor learning is difficult to observe, because it takes place every hour of the awake and behaving time when infants work on they motor skills with intensity and dedication.

The development of reaching is also hard to study theoretically. First, there is currently no complete neurally grounded theory of reaching! A model from our own group (Martin, Scholz, & Schöner, 2009) integrates process accounts for the generation of timed movement commands, for movement initiation and termination, for solving the degrees of freedom problem, and for muscular control. An account for the scene and object perception, and the processes through which such perceptual information is transformed into motor frames is missing. Also, the solution of the degree of freedom problem is not neurally grounded and questions persist about how well the muscle model solves control problems in reaching. A model based on the framework of optimal control (Shadmehr & Krakauer, 2008) does a better job at addressing these control problems, but is far from neurally grounded, does not address muscle properties, and

is also missing a perceptual and movement preparation component. A recent model of learning to reach (Caligiore et al., 2014) takes muscle properties into account while providing an account for control, and addresses the neural grounding of movement parameters, but is not addressing movement initiation and termination nor scene perception and the associated coordinate transformations. The most neurally grounded models of reaching are still those from the Grossberg school (Cisek, Grossberg, & Bullock, 1998), which have been mapped in quantitative detail to neural data for some components (Cisek, 2006). They fall short of addressing scene and object perception as well as the initiation and termination of movements. Closest to integrating all processes come recent proposals based on Dynamic Field Theory (Strauss & Heinke, 2012; Strauss, Woodgate, Sami, & Heinke, 2015; Fard, Hollensen, Heinke, & Trappenberg, 2015). The synthesis (Zibner, Tekülve, & Schöner, 2015a, 2015b; Tekülve, Zibner, & Schöner, 2016) we review here is very similar in spirit and strongly overlaps with these proposals with respect to scene perception and motor planning. We elaborate movement timing, motor control, and the sequential organization of movement in more detailed ways, using different process models.

Second, understanding how reaching develops requires an understanding not just of motor learning, but of learning in the many different component processes from perception to control. Current accounts typically focus on an individual component. The most complete model so far (Caligiore et al., 2014), for instance, focuses on learning to control the arm by predicting the torque profiles needed to reach. Other models look at the learning of kinematic movement plans (Schlesinger, Parisi, & Langer, 2000), of kinematic models (Sun & Scassellati, 2005; Herbort & Butz, 2009; Narioka & Steil, 2015), or of feedback control parameters (Savastano & Nolfi, 2013).

Finally, accounts for learning are often based on a learning regime in which the organization in time of the learning process is not achieved autonomously by the learning system itself, but imposed from the outside (see Sandamirskaya and Storck (2015) for discussion). In particular, autonomously learning from experience requires much processing “infra-structure”. For instance, to learn from the correlation between a motor command and the perceived outcome, as assumed in many of the “motor babbling” accounts, the movement system must be able to keep the motor command in working memory to bridge the considerable delay between the two neural events.

Goals of this chapter

In this chapter our goal is to review the component processes that must be in place to successfully reach for and grasp objects. The review will be conceptual first, but will also discuss relevant neural principles. We will illustrate the concepts through a concrete neural dynamic model that provides a process account from sensory inputs to generating movement of a simulated biomechanical plant. The model is not quantitatively anchored in experimental data, but we do discuss the link to experimental signatures.

A second step we want to make is expose what happens when some of the components are not in place. By comparing the various ways the model may break down

in a form of reverse development, we aim to uncover potential behavioral signatures of development, and also highlight all the problems the human nervous system must solve to reach successfully.

A key contribution of the model is that the reaching behavior itself is entirely autonomous. There is no hidden “script” that activates each component process as required (as is the case for many robotic demonstrations of reaching). In particular, we will show how correction movements emerge from the time continuous neural dynamics when an initial reach is not complete, and use this observation to account for the emergence of multiple movement units.

Finally, within a limited setting we will show how the organization of motor acts can be learned autonomously from behavior. This demonstration of principle will highlight all the processing infra-structure required to achieve such autonomous learning and provides a perspective for what a full account for learning to reach may need to address.

2 A neural process account of reaching

In a neural process account of reaching, the neural networks of the brain are linked to sensory and motor systems to bring about the motor behaviors that achieve the reaching act. Such an account may be contrasted to abstract “curve fitting” models, which by themselves do not explain how behavior emerges. For instance, to postulate that an infant selects the interpretation of a stimulus that maximizes its likelihood does not explain how its nervous system actually does that even while such a model may provide a fit to certain psychometric curves.

Theoretical models that are process accounts may ultimately be linked to real sensors and real actuators, in robotic demonstrations of reaching. Of course, such robotic demonstrations are never perfectly consistent with what is known about how organisms generate movement. No actuator perfectly mimics real muscles nor do technical sensors mimic the neural function of the retina or other sensory organs.

A neural process account is always based on the choice of a particular level of description that entails a particular level of neurally mechanistic detail. That choice is constrained by what is known about the neural substrate of motor behavior. In spite of enormous progress, the exact neural circuitry underlying motor behavior remains unknown (Lisman, 2015). Neural process accounts are thus usefully founded on neural principles, rather than on a detailed description of the neural substrate and associated neural mechanisms. We know, of course, that the central nervous system is not a digital computer and does not have a CPU, into which it can load data to operate on. Thus, process accounts that depend conceptually on algorithms are not by themselves compatible with neural principles.

The most commonly accepted level of description at which neural principles can be articulated and used to develop process accounts of behavior is the level of neural population activity (Georgopoulos, 1986; Bullock, Cisek, & Grossberg, 1998; Cisek & Kalaska, 2005). The framework of Dynamic Field Theory (DFT) is positioned at this level (Erlhagen & Schöner, 2002; Schöner, Spencer, & the DFT Research Group, 2015).

In DFT, properties of the strongly recurrent neural networks that generate behavior are formalized into a set of mathematically expressed concepts. Central is the notion of stability, the capacity of neural activation patterns to resist change. Rather than review yet again the neural principles formalized in DFT, we introduce them as we go through the neural process model of reaching behavior and its development.

What is entailed in making a reaching movement oriented toward an object? Figure 1 illustrates the five processes that are minimally required to generate such behavior.

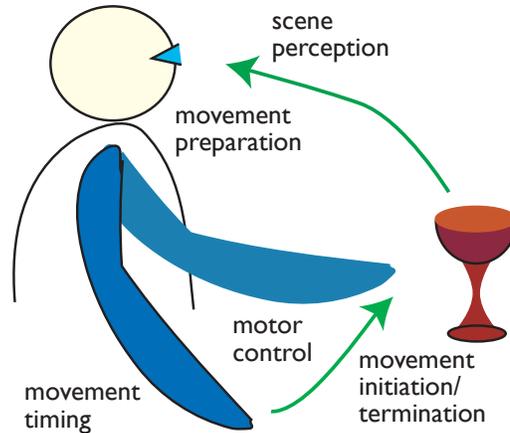


Figure 1: A schematic illustration of five component processes entailed in generating a reaching movement aimed at an object.

2.1 Scene perception

Object oriented movement behavior requires, first of all, perception of the environment and attentional selection of the object, toward which the behavior is directed. Such scene perception is predominantly visually driven, although reaching toward haptically identified or memorized object locations is possible. Humans are very good at perceiving and memorizing visual scenes, much better than at memorizing arbitrary material. For instance, individuals who looked at 10 objects per naturalistic scene, for 10 seconds each, had 90% recall of object identity or pose even a week later (Hollingworth, 2004). Our visual cognition is particularly well tuned to this problem.

Scene perception clearly is tightly linked to looking and attention. Only portions of a scene that have been attended to are memorized well enough to detect change (when transients are masked, (Simons, 2000)). The attentional selection of visual objects as well as the selection of gaze locations is commonly accounted for through the notion of visual salience that characterizes the stimulus properties determining the probability of attraction gaze or covert attention (Itti & Koch, 2000). The process of such selection decisions is captured by strongly recurrent neural networks, formalized as neural dynamics (Kopecz & Schöner, 1995; Schöner et al., 2015), as illustrated in Figure 2. In the neural field version of neural dynamics, visual space is represented

by an activation field in which localized peaks of activation indicate a selected visual location. Such peaks may be induced by localized input, which may reflect the saliency of the visual array, and are stabilized against decay by local excitatory interaction. Selection is enabled by longer-range inhibitory interaction. When neural interaction within the activation field is sufficiently strong, peaks may be sustained even as localized input is removed. Sustained activation provides an account for visual working memory (Johnson, Spencer, Luck, & Schöner, 2009).

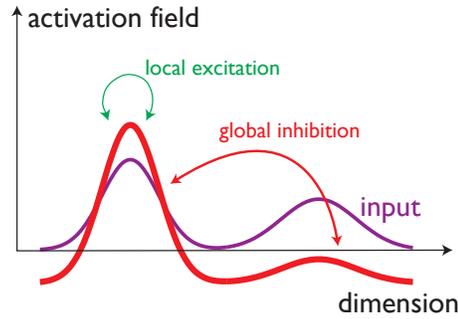


Figure 2: A neural activation field is defined over some dimension and receives bimodal input. The field selects the left-most local maximum of input through local excitatory interaction, which lifts activation to above threshold values at that location (more activation than input) and global inhibition which suppresses activation at other locations (below the level of input).

Here is a very brief mathematical tutorial on dynamic neural fields, expressed for a neural activation field, $u(x, t)$, defined over a single dimension, x (in the model, the fields comprise multiple dimensions such as the two-dimensional visual array or two movement parameters). Activation evolves according to this integro-differential equation:

$$\tau \dot{u}(x, t) = -u(x, t) + h + s(x, t) + \int dx' w(x - x') g(u(x, t)). \quad (1)$$

The terms up to the integral are a time continuous version of an input driven neural dynamics as it is commonly used to model the time courses of neural activation. Without input, s , activation is in a resting state, $h < 0$, that is stable with time constant, τ . Weak localized input, $s(x)$, shifts that stable state to $u(x) = h + s(x)$. The neural dynamics (through the “ $-u$ ” term) creates time courses through which activation tracks changes in input $s(x, t)$ as an exponential low-pass filter.

This input driven regime becomes unstable once the activation level at any location, x , approaches the threshold at zero defined by the sigmoidal nonlinearity, $g(u)$ (in the modeling below chosen as $g(u) = (1 + \beta u / (1 + \beta |u|)) / 2$). This is when neural interaction within the field engages as $g(u)$ becomes positive. Locations, x' , close to the activated location, x , provide excitatory input ($w(x - x') > 0$). Locations, x' , far from the activation location, x , provide inhibitory input ($w(x - x') < 0$). This pattern of neural interaction within the field stabilizes peaks of activation centered on the location at which input first pushed activation beyond threshold. When activation is pushed above threshold at more than one location, such peaks may instantiate selection decisions in which one activation peak suppresses peaks at alternate locations. Which location is selected depends on the timing and strength of input, on the prior patterns of activation, and on random fluctuations of activation. When neural interaction is sufficiently strong, peaks may be sustained even after the inducing localized input is removed. The location of the peak then encodes a working memory for the past selection decision.

Because scene perception involves gaze and attentional shifts, only a part of the visual array is in the attentional foreground at any moment in time. Scene perception is thus largely based on memory (Hollingworth, 2004), in what may be more appropriately called scene representation.

Early forms of reaching are not always associated with looking at the object (von Hofsten, 1984). As reaching develops, that link becomes closer. Coordinating attention and reaching is, therefore, a developmental achievement rather than a logical necessity. After reaching has been established, the relationship between looking and reaching remains complex, however. Corbetta and associates (Corbetta, Thurman, Wiener, Guan, & Williams, 2014) presented infants with object large enough for them to reach toward different locations on the object. These infants were followed longitudinally from about 2 to 12 months of age. Early in that developmental window, the location

at which infants looked did not predict the location toward which infants reached. Over time, infants looked more to the location on the object to which they reached.

In this chapter, we will use a simplified neural process model of reaching to illustrate ideas. Figure 3 provides an overview over the architecture. The entire topic of scene perception has been trivialized in this model by assuming that a distribution of activation defined in body-centered coordinates is available to the processes of movement preparation. Visual locations of reachable objects are marked by localized maxima of that distribution of activation.

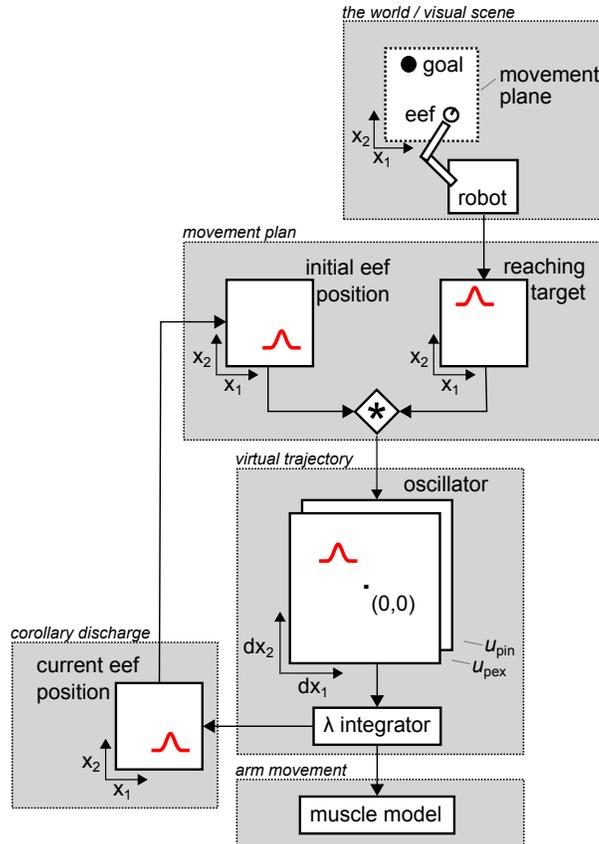


Figure 3: A survey over the neural dynamic architecture used in this chapter to illustrate the five key processes of movement generation

2.2 Movement preparation

There is ample behavioral and neural evidence, that movements are prepared ahead of their initiation (see Erlhagen and Schöner (2002) for review). Thus, for instance, the time needed to initiate movement after a stimulus has specified the movement goal, the reaction time, reflects the metrics of the movement alternatives. If those alternatives are metrically closer to each other, reaction time is shorter, reflecting more overlap between the neural activation states that correspond to either movement (McDowell,

Jeka, Schöner, & Hatfield, 2002). In fact, these neural activation states can be directly observed in motor and pre-motor cortex in the form of peaks of activation in a distribution of population activation (Georgopoulos, 1986; Bastian, Schöner, & Riehle, 2003; Cisek & Kalaska, 2005). Activation even in motor cortex precedes movement initiation and predicts movement parameters. Finally, at the kinematic level, reaching movements directed at an object start out in the direction of the object in adults, so that from the first milliseconds of its trajectory, the path of the hand and the movement time can be predicted. This is a consequence of the robust kinematic regularity that characterizes adult movement (Soechting & Lacquaniti, 1981).

It is the movement parameters, that characterize movements as a whole, that have specific values from the very start of the motor act. Most prominent among these is the direction of the hand’s movement in space, the extend of the movement of the hand through space (amplitude), the overall duration of the movement (movement time) and other parameters such as the anticipated level of resistance to the movement. Movement preparation thus means determining the values of such movement parameters.

Once an object has been selected in the scene representation as the target of a reaching movement, specifying kinematic movement parameters like movement direction and amplitude involves some simple geometrical computations (Bullock & Grossberg, 1988). The direction, for instance, is the angle that the line connecting the initial position of the hand to the object forms relative to some reference axis (see Figure 4) . Such computations are trivial to implement on a computer, but not in a neural network. Neural networks do not take an “argument” and “operate” on it. They need a particular pattern of connectivity that brings about the computation and that is linked to a neural representation of the argument. In this instance, the computation amounts to a coordinate transform (Figure 4): If the spatial representation of the target is transformed into a reference frame that is centered in the initial position of the hand, then the direction and amplitude of the reaching movement can be read off with a fixed neural mapping. For instance, all field locations along a line from the center out vote for the corresponding movement direction. Similarly, all field locations on a circle of a given radius vote for the matching movement amplitude.

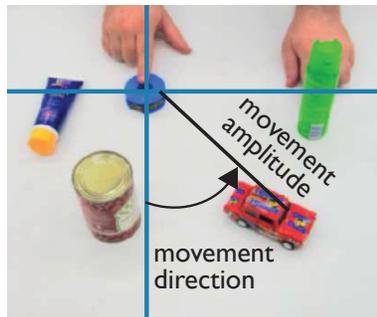


Figure 4: The kinematic movement parameters that characterize a reach toward an object may be obtained by transforming the neural representation of the location of the object into a frame of reference anchored in the initial position of the hand.

The mathematics of the coordinate transform from a visual reference frame to a frame anchored in the initial position of the hand is sketched here. Let $u_{\text{ini}}(\mathbf{x}, t)$ be a neural field that represents the initial position of the hand through a peak at the appropriate location (\mathbf{x} is a two-dimensional position vector that spans all possible initial positions of the hand). The location of the target object is represented as a peak in the neural field, $u_{\text{tar}}(\mathbf{x}, t)$, defined over the same two-dimensional space (both spatial representations could be thought of being body-centered). A hand-centered neural representation of the target object can be obtained from this neural map:

$$s_{\text{plan}}(\mathbf{x}, t) = \iint dx'_1 dx'_2 \sigma(u_{\text{ini}}(\mathbf{x} - \mathbf{x}')) \sigma(u_{\text{tar}}(\mathbf{x}', t)) \quad (2)$$

Every location, $\mathbf{x}' = (x'_1, x'_2)$, in the target field has a connection to every location, \mathbf{x} in the movement planning field. At any given moment, only a subset of these connections is active, those selected by the peak in the field that represents the initial position of the hand. The selection mechanism formalized here is a form of “shunting” in which the activation in one set of neurons can turn of and on the connection between two other sets of neurons (the map is also referred to as “sigma-pi” connectivity). This is not the only possible neural implementation of the transform (see Pouget and Snyder (2000); Schneegans (2015) for more). The transform is essentially a “steerable” neural map: the mapping from one space, in which the visual object is represented, onto another space, in which the movement direction is represented, is steered by the initial position of the hand. Such transforms are costly: many connections and many neurons are needed and their connectivity pattern needs to be just right. The maps create invariance: no matter where the hand is initially posted, the direction of the movement can be determined. The learning of such transforms is a likely developmental challenge (Sandamirskaya & Storck, 2015).

2.3 Movement timing

Reaching movements, like most voluntary human movement, are timed. That is, the hand’s trajectory and velocity profile has a characteristic shape that is reproducible across repetitions and scales with movement amplitude (Morasso, 1981). If a movement sequence (for example, in hand writing) is performed at different spatial scales, the relative timing of the corresponding pieces of movement (the arc forming the letter “e”, for instance), remains invariant (the “e” takes up the same percentage of the total movement time) (Viviani & Flash, 1995). Reaching movements are coordinated across limbs so that when timing demands on one limb are varied, the movement timing of the other limb adjust to retain the same relative timing (Kelso, Southard, & Goodman, 1979). Reaching movements may also be coordinated with perceived events such as when moving objects are intercepted. That coordination is stable in the sense that the reaching movement is updated in response to visual information about the timing of

the movement target (Brenner, Driesen, & Smeets, 2014). Coordination is maintained even where it is not strictly necessary. The opening and closing movement that achieves grasping is coordinated with the transport component that moves the hand toward the grasping object (Jeannerod, 1984). In principle, the hand could just open and stay open until the object is reached. Coordination means, however, that a later reach is accompanied by a later opening of the hand for the grasp as well.

The theoretical understanding of timing is based on the notion of an oscillator (Schöner, 2002). More precisely, clocks are dynamical systems that generate reproducible and stable time courses formalized as stable limit cycles. The stability of limit cycles means that the time courses generated resist change. Coupling multiple limit cycles generically leads to their synchronization, an account for coordination.

The concept of a dynamical system generating stable time courses is not restricted to accounts of periodic movement. Temporally discrete motor acts such as a single reaching movement can be understood on the same basis (Schöner, 1990). The idea is that the dynamical system generates a stable time course, which ends, however, in a stable stationary state. (We will address below the processes of initiation and termination required in such a view.) More specifically, a neural oscillator consists of an excitatory population that is reciprocally coupled to an inhibitory population (Amari, 1977; Ermentrout, 1998). An initial stimulus sets in motion an “active transient” in which activation first rises, then falls as inhibition cancels activation. Different dynamic properties of such neural oscillators may generate trajectories of differing amplitude and duration. Our account of reaching movement postulates an ensemble of such neural oscillator. When a peak forms that represents a movement plan, it drives a subpopulation of these neural oscillators that has the appropriate amplitude and movement time to ultimately generate a successful reaching movement (after going through kinematic transformations, muscle activation, and biomechanics, to be discussed below). The projection from these neural oscillators to the downstream neural processes is learned so as to achieve these movement goals (Figure 3).

There are good arguments as to why the timing signals are generated in spatial terms, such as the direction and speed of the hand’s movement in space. Neural data are supportive of that idea (Schwartz, 1994; Moran & Schwartz, 1999) as are data about movement coordination (Mechsner, Kerzel, Knoblich, & Prinz, 2001) that are consistent with the ideomotor principle according to which movement is generated in the same reference frame in which it is perceived. The ease with which we coordinate the hands timing with the motion of perceived objects is suggestive of such a reference frame as is the invariance of movement timing with scale (see above).

That movement timing poses a developmental challenge is intuitive, although there seems to be little direct empirical evidence about the development of movement coordination. Counter-intuitively, analysis of spontaneous kicking movements in infants revealed that in-phase and alternating coordination patterns are observable from earliest infancy (Thelen, 1981).

A possible mathematical form of the neural dynamics of movement timing invokes two fields that together form neural oscillators. The excitatory layer, $u_{\text{ex}}(\mathbf{x}, t)$, and the inhibitory layer, $u_{\text{in}}(\mathbf{x}, t)$ are coupled as follows:

$$\begin{aligned} \tau_{\text{ex}} \dot{u}_{\text{ex}}(\mathbf{x}, t) &= -u_{\text{ex}}(\mathbf{x}, t) + h + s_{\text{plan}}(\mathbf{x}, t) + c_{\text{mov}} g(u_{\text{mov}}^{\text{int}}(t)) \\ &\quad - \int d\mathbf{x}' w_{\text{ex, in}}(\mathbf{x} - \mathbf{x}') \rho(u_{\text{in}})(\mathbf{x}', t) \end{aligned} \quad (3)$$

$$\tau_{\text{in}} \dot{u}_{\text{in}}(\mathbf{x}, t) = -u_{\text{in}}(\mathbf{x}, t) + h + s_{\text{plan}}(\mathbf{x}, t) + c_{\text{mov}} g(u_{\text{mov}}^{\text{int}}(t)) \quad (4)$$

where $\tau_{\text{ex}} < \tau_{\text{in}}$ and $w_{\text{ex, in}}(\mathbf{x} - \mathbf{x}')$ is an interaction kernel with local excitatory and lateral inhibitory interaction. Anywhere across these homogeneous fields, a local region may generate an active transient, once it receives input from the motor plan, $s_{\text{plan}}(\mathbf{x}, t)$. Activation in both excitatory and inhibitory layers then rises. The slower rise of inhibition begins to suppress activation in the excitatory layer, eventually suppressing activation below the threshold of the semi-linear transfer function

$$\rho(\mathbf{x}, t) = \begin{cases} u(\mathbf{x}, t) & \text{for } u(\mathbf{x}, t) > 0 \\ 0 & \text{else.} \end{cases} \quad (5)$$

This dynamics thus performs a one-shot active transient in response to input. Input by a neural node, $u_{\text{mov}}^{\text{int}}$, is required to induce the active transient. This node is part of a neural dynamics of behavioral organization, that generates the sequence of initiation and termination events that organizes reaching movements (see Section 2.5).

The field of neural oscillator generates a (virtual) velocity vector

$$\mathbf{v}(t) = \iint dx_1 dx_2 \omega(\mathbf{x}) \rho(u_{\text{ex}}(\mathbf{x}, t)). \quad (6)$$

as its output. The weight function, $\omega(\cdot)$, is learned using gradient descent as a linear function of $\mathbf{x} = (x_1, x_2)$. These weights define tuning curves of the neural oscillator field, in which each field site is sensitive to both the orientation and the distance (and thus peak velocity) of the movement.

2.4 Motor control

To physically generate the movement, actuators must be given motor commands in such a way that deviations from average trajectory are limited by stability properties of the plant or the control. In human reaching movements, the actuators are, of course, muscles attached to the skeletal musculature. The fact that limbs are actuated at the joint level alone generates challenges for motor control: As a joint is accelerated, interaction torques act on other joints that may move these off a commanded path. For instance, if you rapidly move your elbow while relaxing your wrist, your hand flaps around. This movement at the wrist is a consequence of interaction torques. The wrist

is in an accelerating reference frame, inducing inertial interaction torques, that also rotates, inducing coriolis and centrifugal interaction torques.

Conventional robot arms deal with such interaction torques with “brute force”: servo controllers keep each joints velocity and position close to the commanded values, applying higher forces to control away these disturbances. Humans motor control is very different. Muscles are relatively weak. As elastic elements, they are generally weaker than the inverse pendulum that body segments form with the vertical gravitational form (Loram & Lakie, 2002). Muscles do have feedback control loops, including those through the spinal cord. But the delay in these reflex loops, on the order of about 50 ms, are relatively long compared to movement times as short as 300 ms.

In what form do motor commands ultimately bring about muscle activation? Because the activation of muscles depends on their length and inner state, motor commands do not directly specify muscle forces. More appropriately one may conceive of muscle commands as parameters modulating the peripheral reflex loops that establish the length dependency of muscle activation (Feldman, 1986, 2011). Specifically, the descending motor command could be viewed as setting a threshold for the activation of the muscle. When the muscle length is above that threshold, muscle activation is induced that leads to the generation of muscle force that contract the muscle. That relationship between muscle length and muscle force implies that muscles are effectively non-linear springs, whose the resting length can be varied by the descending motor command.

Because motor commands are thus spatial in nature, they may be derived from the spatial timing signal through a kinematic transform: From the hand’s movement direction and speed to muscle lengths and their rates of change. This transformation is complex and contains problems not fully understood at this time. First, multiple muscles act on each joint. The muscles that act in opposing directions of movement, agonists and antagonists, may have different levels of co-contraction that do not affect the joint angle, but joint stiffness. Joint stiffness is not specified by the kinematic transformation and the control of the related joint compliance is one of many open issues. In our model, we assume a particular level of stiffness that remains constant through movement, and thus neglect the known variation of joint stiffness during movement. Second, the redundancy of muscles acting in the same direction on a joint is also a topic of much research. We neglect that issue by lumping all muscles acting on a joint into a single muscle model. The muscle lengths are thus combined into a single joint angle. This is a “virtual” joint angle that describes the motor command given to this ensemble of muscles. The real joint angle then evolves in a way that may not always be close to the virtual joint angle depending on interaction torques and other factors. Third, the kinematic transform faces kinematic redundancy: In most movement systems and certainly for reaching, there are more joints available to achieve the movement of the relevant end-effector (here the hand in space) than necessary. There is extensive research on how these many degrees of freedom are harnessed to realize a particular trajectory of the hand in space (Todorov & Jordan, 2002; Latash, Scholz, & Schönner, 2007; Martin et al., 2009). We condense all that into a kinematic transform that at any moment in time selects one particular virtual joint angle velocity vector

across all joints.

The mathematical formalization of the sketched ideas is as follows: The virtual velocity vector, \mathbf{v} , is transformed into a virtual joint velocity vector, $\dot{\boldsymbol{\lambda}}$, using the pseudo-inverse (upper index $+$) of the manipulator Jacobian, \mathbf{J} , that depends on the current virtual joint configuration $\boldsymbol{\lambda}(t)$:

$$\dot{\boldsymbol{\lambda}} = \mathbf{J}^+(\boldsymbol{\lambda}(t)) \mathbf{v}(t). \quad (7)$$

The virtual joint trajectory, $\boldsymbol{\lambda}(t)$, is obtained by integrating over time. The virtual joint trajectory defines a set of motor commands to all muscles that transitions from an initial to a final equilibrium length of each muscle. In a simplified, second order linear model of the muscle and its reflex loops, the virtual joint trajectory sets the time-varying equilibrium point for each joint, combining the contributions of all agonist and antagonist muscles converging on each joint:

$$\ddot{\boldsymbol{\theta}} = -\mathbf{K}(\boldsymbol{\theta} - \boldsymbol{\lambda}) - \mathbf{B}\dot{\boldsymbol{\theta}}, \quad (8)$$

where $\boldsymbol{\theta}$ is the joint angle vector, and \mathbf{K} and \mathbf{B} are the effective stiffness and viscosity matrices of the muscle groups.

The simple model does not address biomechanics. The kinematics of a seven degrees-of-freedom arm was solved numerically to compute the hand’s trajectory in space (see Section 3).

2.5 Movement initiation and termination

All voluntary movement consists of sequence of events. A reach initiates and ends, but may also entail the initiation of the opening movement of the hand and then its closing for a grasp. Other movements may follow once an object has been grasped. A reaching movement may be interrupted or aborted at any time (Bullock & Grossberg, 1988).

The sequential organization of movement is a special case of the problem we like to call *behavioral organization*, a problem much overlooked in neural models of behavior and cognition (Graziano, 2006; Sandamirskaya, 2015; Pezzulo & Cisek, 2016). In motor control, models will often focus on a single phase of the motor act, such as the phase of movement generation itself. Cognitive models may talk about the serial order of events, but fail to address how these events are organized in time and interact with the time course of behavior generation between events (Botvinick & Plaut, 2004).

An account for behavioral organization is necessary to endow neural models with a form of autonomy: No outside “script” or algorithm must be invoked to understand how neural systems move from one task to another. To date, many connectionist models are limited in their degree of process autonomy. This is true for learning processes as well, in which a learning protocol is often imposed from outside of the connectionist model (we will come back to this topic in Section 4).

One particular challenge to a neural process account for behavioral organization is the need to understand the relationship between events that occur at discrete moments in time, the transitions from one action to another, and the movements themselves that take varying amounts of time. For instance, when a reaching movement takes longer because an obstacle is avoided (Grimme, Lipinski, & Schöner, 2012), then the switch to the next phase of the object-oriented action must occur later.

Clearly, behavioral organization is a developmental challenge. Infants do not get the orders of movement tasks right a lot of the time (von Hofsten, 1984). And getting them right is a prerequisite to success.

In saccadic eye movement, the neural basis of aspects of behavioral organization is understood somewhat. A population of neurons in the superior colliculus is known to be active, while the animal fixates gaze on a visual structure. A saccade is initiated when the fixation population becomes deactivated. Theoretical accounts for this form of transition (Kopecz, 1995; Wilimzig, Schneider, & Schöner, 2006) were the origin of the ideas we now sketch (Richter, Sandamirskaya, & Schöner, 2012). The key idea is that motor behaviors are activated or deactivated by dedicated neural nodes that reflect the activation in small populations of neurons that have common input. These nodes may shunt neural connections or act by boosting or suppressing activation in a population. Only neural populations in the self-stabilized state of supra-threshold activation are “active” in the sense that they may can impact on the downstream processes. The transition from inactive to active takes the form of the detection instability discussed earlier. The transition from active to inactive takes place of through a separate, reverse detection instability. In both cases, even graded changes of input lead to events at a discrete moment in time, when the instability is induced.

The neural nodes modulate the conditions at which these instabilities may occur. They are themselves governed by the same type of neural dynamics and thus their own activation may undergo the dynamic transition from “off” to “on” and vice versa. Three types of neural nodes are distinguished based on their functional role in organizing sequences of neural states. “Intention” nodes represent neural populations that enable a behavior by boosting its processes to push them through the detection instability if the appropriate sensory inputs are available. “Condition of satisfaction” (CoS) nodes become activated when input from an associated intention node is matched by input from sensory or internal sources predicted to signal the successful completion of the intended motor behavior (Sandamirskaya & Schöner, 2010). “Condition of dissatisfaction” (CoD) nodes becomes activated when input from an associated intention node is matched by input from sensory or internal sources predicted to signal the failure of the intended motor behavior.

Mathematically, the sequential organization of movement is achieved by a neural dynamic architecture illustrated in Figure 5 (and expanded in Section 4 to include grasping and looking behaviors). While the intention node, $u_{\text{reach}}^{\text{int}}$, is active, reaching is being pursued. A single timed movement act is initiated when the intention node, $u_{\text{mov}}^{\text{int}}$, becomes active. The neural dynamics of this node

$$\begin{aligned} \tau \dot{u}_{\text{mov}}^{\text{int}}(t) = & - u_{\text{mov}}^{\text{int}}(t) + h + c_{u,u} g(u_{\text{mov}}^{\text{int}}(t)) + c_{\text{mov,reach}} g(u_{\text{reach}}^{\text{int}}(t)) \\ & - c_{\text{int,CoS}} g(u_{\text{mov}}^{\text{CoS}}(t)) - c_{\text{int,CoD}} g(u_{\text{mov}}^{\text{CoD}}(t)). \end{aligned} \quad (9)$$

is analogous to that of activation fields (Equation 1), with the interaction kernel being replaced by self-excitation of strength, $c_{u,u}$. The *activated* state of the node with an activation level above the threshold of the sigmoid function, $g(\cdot)$ is analogous to the self-stabilized peaks of activation in neural fields. The node becomes activated when input from the reaching intention node, $u_{\text{reach}}^{\text{int}}$, pushes the sub-threshold state, $u_{\text{mov}}^{\text{int}} \approx h + \text{inputs}$, to the detection instability. Inhibitory coupling from Condition of Satisfaction (CoS) and Condition of Dissatisfaction (CoD) nodes may induce loss of stability of the activated state in the reverse detection instability.

Once the intention to move is activated, it enables the neural timing field to generate a virtual velocity trajectory: In Equation 4, an activated $u_{\text{mov}}^{\text{int}}$ is necessary for localized input from the movement plan, $s_{\text{plan}}(\mathbf{x}, t)$, to drive the neural oscillator field above threshold to generate an active transient (somewhat analogously to the “go” signal of Bullock et al. (1998), although it does not shape the velocity profile).

The end of a single time movement is signaled by the activation of the condition of satisfaction, $u_{\text{mov}}^{\text{CoS}}$:

$$\begin{aligned} \tau \dot{u}_{\text{mov}}^{\text{CoS}}(t) = & - u_{\text{CoS}}^{\text{int}}(t) + h + c_{u,u} g(u_{\text{mov}}^{\text{CoS}}(t)) \\ & - c_{\text{CoS,ex}} g(u_{\text{ex}}^{\text{pd}}(t)) + c_{\text{CoS,in}} g(u_{\text{in}}^{\text{pd}}(t)) \end{aligned} \quad (10)$$

that is driven by the oscillator field through peak-detectors (upper index pd) that spatially integrate over the output of an activation field. A peak in the inhibitory layer excites, a peak in the excitatory layer inhibits. Thus, the CoS is activated one a peak in the inhibitory, but not in the excitatory layer, signals that a single movement event has taken place.

The intention to move may also be deactivated if the neural representations of the initial position of the hand or the target are lost (Figure 5), signaled by the CoD. The position of the hand is continuously updated by integrating in time the virtual velocity. This corollary discharge signal updates the neural representation of the initial position of the hand when the movement has been terminated. The reaching intention remains activated until a match is detected between the updated initial hand position and the target position, so that the hand has reached the target.

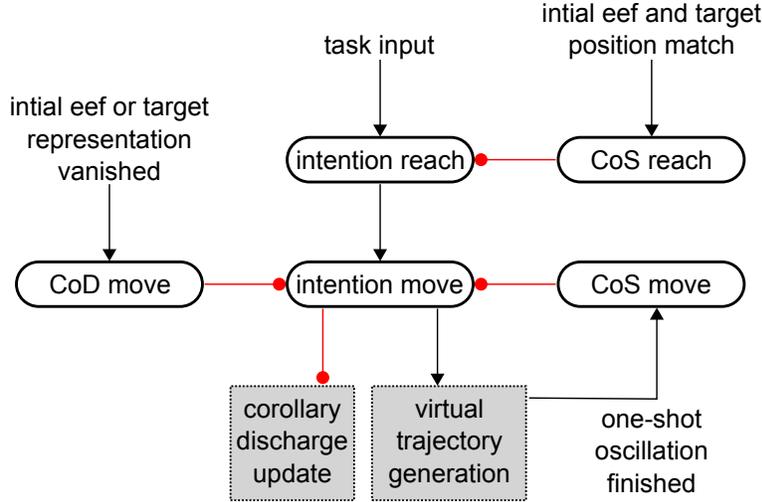


Figure 5: A part of the neural dynamic network that organizes reaching movements. Arrows indicate excitatory input, lines with a circle at its end indicated inhibitory input.

3 Results

To illustrate properties of the neural process model of movement generation and its developmental implications we show simulation results of the complete model. For all experiments, we use artificial visual inputs in the form of localized peaks of activation in a simulated visual array. Compared to using real camera input, this gave us full control of stimulus strength and position for reproducibility. We used the simulation solution Webots (<http://www.cyberbotics.com>) to execute the movements with the seven degrees-of-freedom arm. This made it possible to simulate neural dynamics that reflected imperfect control (the “impaired” system, see Section 3.2) and produced jerky motor commands without risking damage to robotic hardware. The correctly tuned architecture was tested on hardware as well (using a RGB camera and the Kuka lightweight arm), but this will not be discussed here (see Zibner et al. (2015a) for more details).

3.1 Making a movement

Figure 6 illustrates how the model makes individual movements from an initial hand position to three different visual targets. The parameters of the model were tuned to put the different components of the model into the dynamic regime described. The movement paths (left) are approximately straight with slight outward curvature. They are generated by the virtual hand velocity profile (right panel, left-most curves). Because movement time was set to be equal across the three movements, peak velocity scaled with movement amplitude. The resulting real velocity profiles of the hand’s movement in space (right panel, right-most curves) are considerably delayed over the

virtual velocity profiles due to the low-pass filter properties of the muscle model. Near peak real velocity, the virtual velocity has already returned to zero. This is consistent with experimental estimates of the timing relationship between virtual and real trajectories (Ghafouri & Feldman, 2001). Note that in more realistic muscle models there are additional sources of delay such as transmission delays in the reflex loop and the slow build up of muscle force due to calcium dynamics (Gribble, Ostry, Sanguinetti, & Laboissière, 1998). The qualitative picture is, therefore, adequate.

Such offsets between the timing signal generated by the neural oscillator system and the realized movement trajectory raise a range of interesting challenges. In particular, the organization of movement sequences must take into account such delays. Moreover, coordination among different effector systems or with external timing events need to be robust against such timing offsets.

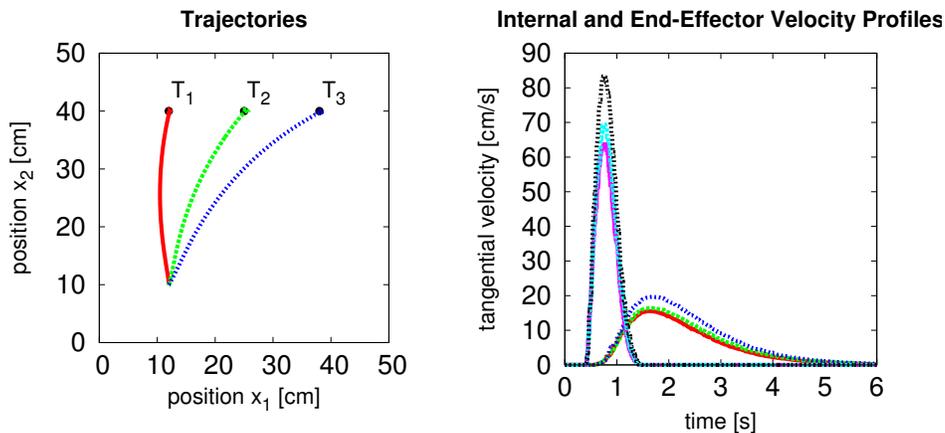


Figure 6: Exemplary trajectories (top left) and profiles of tangential velocity for virtual movements (top right) and end-effector movements (bottom right) for different movement targets. The bottom left plot shows a combination of virtual and external profiles to show that the virtual movement ends roughly at peak velocity of the end-effector movement.

3.2 Developmental signatures

We argued that all components of the movement generation system must be in place and functional in order to bring about reaching movements directed at objects. The simplified integrated model enables us now to examine the ways object-oriented movement fails when different components of the model are not in place. We examine three different ways the model may be “degraded” to explore potential developmental signatures of early stages of reaching. (A) Movement plan: The lateral neural interaction within the neural field that represents the initial position of the hand was weakened. As a result, the working memory of that initial position required for the duration of the movement becomes unstable, leading to a potential loss of the movement plan. The

Table 1: Mean and variance of number of movement units and straightness index for different impairment conditions

Conditions	A	B	C	A+B+C	N
Mean # of MUs	3.00	4.00	1.89	7.56	1.00
Variance	0.00	0.00	0.61	7.53	0.00
Mean straightness	0.998	0.983	0.845	0.941	0.989
Variance	< 0.001	< 0.001	0.017	0.002	< 0.001

virtual movement trajectory may then be terminated prematurely. (B) Movement timing: Neural connectivity from the field of neural oscillators that generates the timing of the virtual trajectory may not be appropriately tuned so that virtual trajectories do not reach the target in a single cycle. This was simulated by setting some of the weights from the oscillator field to random values. (C) Motor control: The kinematic mapping from the virtual hand trajectory to the joint-level motor commands may not be appropriately tuned. This was modelled by adding a random offset to the inverse kinematic mapping.

We performed experiments in each of these three forms of “impairment” as well as when all three forms were combined (A+B+C). In each condition, we simulated movements from three starting positions to three target positions, repeating each simulation once with different random values of different forms of degradation. This made for a total of $2 \cdot 9 = 18$ repetition. We analyzed the resulting trajectories and velocity profiles for the number of movement units (MUs) and for the straightness of the movement path and computed the variance of either measure across the 18 trials. Movement units are segments of the real hand velocity profile between two minima of velocity. Straightness was measured as done in experimental analyses (Thelen et al., 1996). Table 1 provides an overview of the results. Individual sample trajectories from one starting position to three target positions are shown in Figure 7, which may be compared to the sample trajectories of the intact system in Figure 6.

When the movement planning system is impaired (A), the velocity profiles are less smooth, featuring more than one movement unit. This is due to the intermittent interrupt of movement generation when a planning peaks is lost. The system autonomously resets the movement plan and make a new movement attempt. When movement timing is impaired (B), the virtual trajectory generated typically does not reach the target. This number of movement is large and the movement path is less smooth as a result. When the motor control system is impaired (C), movement paths are no longer approximately straight. Multiple movement units are observed as the system does not reach the target in response to a single virtual velocity profile. The combination of all impairments generates variance in the movement path and velocity profile that may begin to approximate reaches of young infants. These observations are born out by the statistical analysis of the two indices analyzed here (Table 1). The developmental path suggests that impairment in motor control induces the strongest variance, followed by impairments in movement timing, followed by impairments in movement planning.

It is actually difficult to generate the kind of macroscopic variance of movement

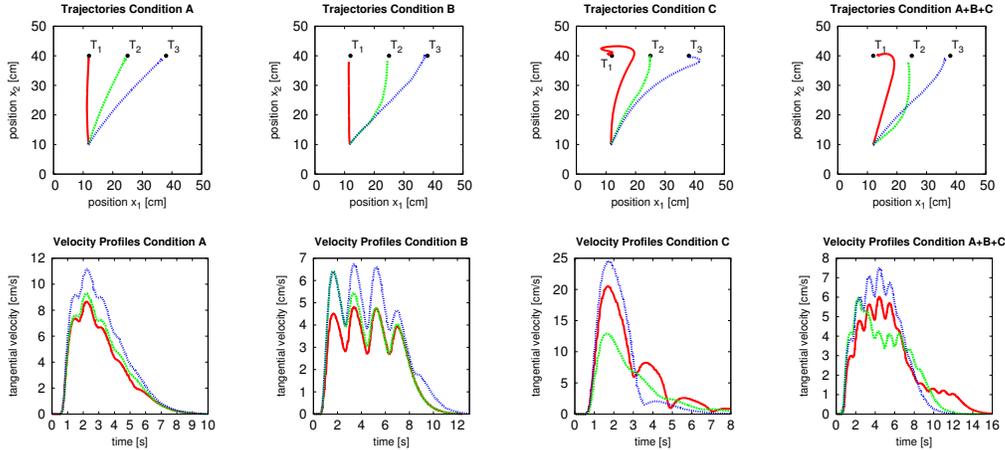


Figure 7: This figure shows trajectories (top) and velocity profiles (bottom) for starting position S_1 to the three targets (solid red for T_1 , bold-dotted green for T_2 , dotted blue for T_3) and, from left to right, the conditions A, B, C, and the combination of all three.

paths and trajectories observed in infants from a neural processing model while still reaching the target. Noise added to the model at the processing level (e.g., in the muscle model) will typically show up as high-frequency fluctuations of the movement trajectories around a mean that is close to the intact model. We believe that the structural sources of variance high-lighted here point the way toward a correct account for movement variance in development.

4 The developmental learning process

We have sketched how different stages of development may be associated with graded differences in the underlying neural processing structure of movement generation. How may changes in these processes emerge from experience? This is the question of autonomous learning: How may learning processes unfold in a spontaneously behaving system rather than in a special “learning mode” in which actions are elicited and rewarded from the outside.

We all know that this is a difficult question. It is difficult in experiment because we cannot readily observe infants and children as they gather experience, practice skills, and learn. Typically, even in longitudinal studies, we may only have a chance to occasionally take a measurement that reflects the current level of competence. But this is also difficult in theoretical modeling. Autonomous learning is not well understood. The majority of neural models are “learned up” in special learning scenarios (Sandamirskaya & Storck, 2015).

We we illustrate the conceptual issues in learning from experience by restricting ourself to one particular subproblem, the problem of learning to sequentially organize different elementary behaviors. We expand the reaching model of the previous section to include looking behavior, opening and closing the hand, as well as moving the hand

back to a resting position. We know that the sequential order of these elementary movement behaviors is not a given as young infants do not necessarily get the order right. von Hofsten (1984) reported that infants of up to four weeks of age generated pre-reaches, movements of the hand in the general direction of an object, even though they often did not visually fixate the object. The frequency of such pre-reaches then decreased while the total time of looking at objects increased. At around 10 weeks of age, pre-reaches re-emerged, now often accompanied by visual fixation of the object. Later, a similar transition led to serial organization of reaching and hand opening and closing.

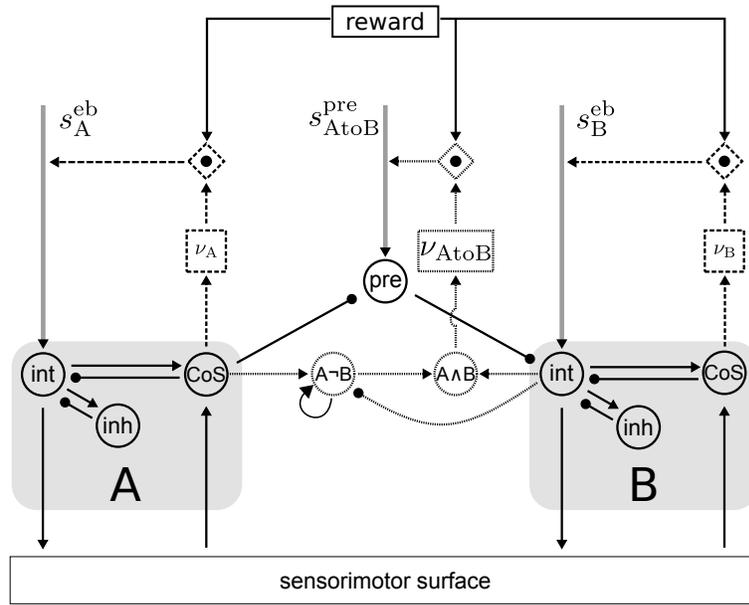


Figure 8: The processing structure for learning the behavioral organization of movement is illustrated for a pair of two elementary movement behaviors, A and B. The same structure exists for all pairs of elementary behaviors. Solid lines are fixed weights assumed to exist from the outset. Dashed lines illustrate weights that are learned autonomously from an initial value of zero. The dashed lines reflect the process structure for learning the boosts of each elementary behavior. The dotted lines reflect processes structure for learning the sequential order of the two elementary behaviors.

Structural support for the sequential organization of different elementary behaviors is a prerequisite for learning that organization. Figure 8 illustrates the neural dynamics for two elementary behaviors, A and B. In the simple model (Tekülve et al., 2016) we use five such elementary movement behaviors: visual fixation, bringing the hand to a resting position, reaching, opening the hand, and closing the hand. Each elementary behavior consists of two neural nodes that directly interact with sensorimotor processes: (1) An intention node, that projects onto the sensorimotor layer and initiates the motor behavior, and (2) a condition of satisfaction (CoS) node that receives input from the sensorimotor layer that signals the completion of the motor behavior. The intention node is self-exciting and coupled to an inhibitory node. Together, these

two nodes may be in different dynamic regimes, including bi-stable “on” and “off” states, a sustained “on” state, and a neural oscillation between “on” and “off”. This is the basis for autonomous activation of an elementary behavior, that plays a role in learning: In a form of “behavioral babbling”, the intention node may activate itself in the presence of noise, and remain activated for a certain amount of time as illustrated in Figure 9.

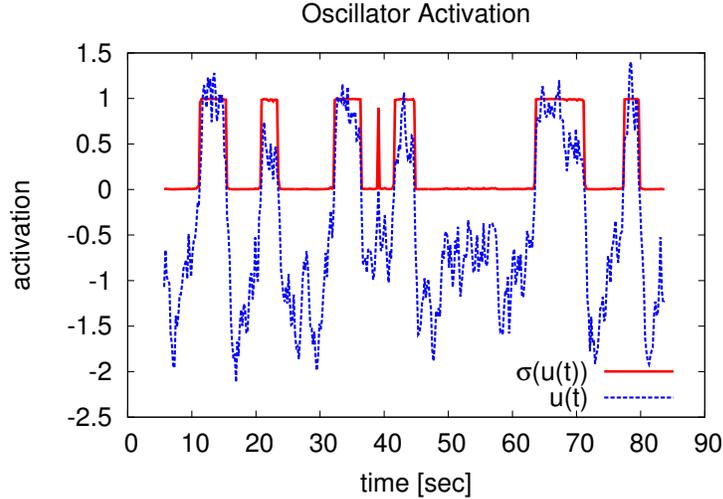


Figure 9: Activation trace of a single intention node (dotted line) and its output (solid line) during behavioral “babbling”. The neural dynamics generates a highly unstable oscillatory pattern.

We also postulate structural support for behavioral organization by assuming that each behaviors condition of satisfaction is coupled to each other behavior’s intention node. This connection is doubly inhibitory, consistent with the organization of behavioral selection in the basal ganglia (Chevalier & Deniau, 1990), through an intermediary “pre-condition” node (Figure 8). If the pre-condition node between behavior A and behavior B is activated, it inhibits its target, behavior B. If the condition of satisfaction of the source behavior, A, then becomes activated, it inhibits the pre-condition node and releases the target from inhibition. This double inhibitory coupling from behavior A to behavior B thus organizes a sequential pattern of activation in which behavior B can only become active once behavior A has been previously activated. The connectivity is effective only if the pre-condition neuron is active in the absence of input from its source.

The learning process varies the resting level of activation of any elementary behavior, s_A^{eb} , and of the pre-condition nodes, s_{AtoB}^{pre} and thus modulates how easy it is to activate which elementary behavior in which order. We were led to introduce yet more specific processing structure for the learning process, inspired by Houk’s model of the basal ganglia (Beiser & Houk, 1998). Two additional interneurons represent different patterns of ordering. One node receives excitatory input from the condition of satisfaction of behavior A and inhibitory input from the intention node of behavior

B. It becomes active, if behavior A was completed, but behavior B was not activated. Through self-excitation keeps that activation in working memory. This node provides excitatory input to a second node which becomes activated once behavior A is activated. That second node thus reflects a particular sequential history of activation, first A, then B.

Autonomous learning faces two conceptual problems. First, activation patterns may come and go. Unless they are actively maintained as working memory (Sandamirskaya & Storck, 2015), these patterns may not necessarily be around when a learning event occurs. This problem is solved, we propose, by keeping memory traces of activation patterns (Erlhagen & Schöner, 2002). These act as low-pass filters, that keep activation patterns available to learning processes for a certain amount of time, that may vary depending on what other, perhaps competing activation patterns arise. Second, autonomously learning from behavior is mediated by neuromodulators. Only at certain moments in time, that are linked to reward, does an “update” of the learning process occur. There is a rich literature on the relevance of timing of reward (Gallistel & Gibbon, 2000) and this problem is central to the mathematical concept of reinforcement learning. In most models of learning, that learning takes place at discrete moments in time is taken for granted and is modeled by using discrete iterative time in the learning rules. Because behavior and neural processes unfold continuously in time, however, this is a problem. One solution lies in the concept of eligibility traces (Sutton & Barto, 1990), in which a population of neurons is transiently activated by each detected “event”, generating a window within which the learning step can take place (Kazerounian, Luciw, Richter, & Sandamirskaya, 2013). Learning is then modulated by the presence of reward and a matching activation pattern.

We provide a brief review of the mathematics of three elements of autonomous learning. (1) The memory trace, $\nu(t)$, reflects the recent history of activation, $u(t)$. The memory trace builds up (on a faster time scale, τ_{build}) whenever the activation, $u(t)$, is above threshold, and decays (on the slower time τ_{mem}) when the activation variable falls below threshold:

$$\begin{aligned} \dot{\nu}(t) &= \frac{1}{\tau_{\text{build}}} \{-\nu(t) + g(u(t))\} g(u(t)) \\ &+ \frac{1}{\tau_{\text{mem}}} \{-\nu(t)(1 - g(u(t)))\}. \end{aligned} \quad (11)$$

(2) An eligibility trace defines the time window in which learning takes place by building a transient reward signal, $r(t)$. In the model, an neural field, u_{rew} , receives input from both the visual perception of the hand and the target. Whenever the two come sufficiently close, a peak is formed in a detection instability. A neural dynamics analogous to the timing model (Equation 4)

$$\tau_{\text{ex}} \dot{u}_{\text{ex}}^{\text{elig}}(t) = -u_{\text{ex}}^{\text{elig}}(t) + h + g(u_{\text{rew}}) - c_{\text{ex,inh}} g(u_{\text{inh}}^{\text{elig}}(t)) \quad (12)$$

$$\tau_{\text{inh}} \dot{u}_{\text{inh}}^{\text{elig}}(t) = -u_{\text{inh}}^{\text{elig}}(t) + h + g(u_{\text{rew}}) \quad (13)$$

(where $\tau_{\text{ex}} < \tau_{\text{inh}}$) translates that discrete movement in time into a transient, time-continuous reward signal $r(t) = g(u_{\text{ex}}^{\text{elig}}(t))$.

(3) The learning rule updates the resting levels of elementary behaviors, s^{eb} , and of pre-condition nodes, $s_{\text{AtoB}}^{\text{pre}}$, whenever their memory traces and the reward signal are simultaneously activated:

$$\dot{s}_{\text{A}}^{\text{eb}}(t) = r(t) \lambda^{\text{eb}} \{\nu_{\text{A}} - (1 - H(\nu_{\text{A}}))\} \quad (14)$$

$$\dot{s}_{\text{AtoB}}^{\text{pre}}(t) = r(t) \lambda^{\text{pre}} \{\nu_{\text{AtoB}} - (1 - H(\nu_{\text{AtoB}}))\} \quad (15)$$

where $H(\cdot)$ is the Heaviside step function, and $\lambda^{\text{eb}} < \lambda^{\text{pre}}$ are the learning rates. Thus, in the presence of a reward signal, only resting levels with matching memory traces are strengthened ($\nu - (1 - H(\nu)) = \nu$), all others decay ($\nu - (1 - H(\nu)) = -1$). In the absence of a reward signal, all resting levels are left unchanged.

We performed simulations in which either of two possible targets were presented at all times (Figure 10). Movements were autonomously generated by behavioral babbling. If the reaching behavior was activated, but now peak in the target field was present because the fixation behavior had not previously been activated, then a peak was induced in the target field at a random location. The hand returned to the resting position when the associated intention node was activated. Paths generated at three phases of the learning process are illustrated. Figure 11 shows how the system first learns to boost the pre-condition node that reflects that visual fixation precedes reaching. This is because trials in which that order occurred where most likely to achieve reward as the visual target impacted on the movement planning. Once this is learned,

the system learns to boost the elementary behaviors fixation and reach. The boost to reaching precedes the boost to visual fixation, because reaching is activated closer in time to the reward and thus has a stronger memory trace when updating occurs.

The movement paths in Figure 10 reflect that pre-reaches first go down in frequency. This is because the newly learned ordering now makes activation of reaches less likely if they are not preceded by visual fixation. Once visual fixation picks up, so does reaching. That pattern matches the observations of von Hofsten.

The opening and closing of the hand is not rewarded in these simulations and therefore, their activation is not promoted during the learning process (Figure 11). This was meant to illustrate that the learning process can autonomously “pick out” the relevant behaviors. In reality, grasping the object may also be a rewarding event, of course.

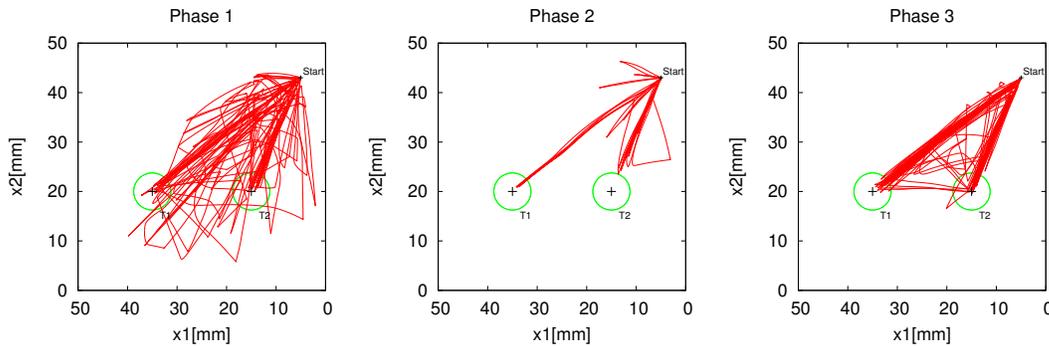


Figure 10: Movement paths generated by the model in three phases of the learning process. Early in learning (left), the reach behavior becomes activated with or without a prior fixation behavior. As a result, movements may go to either the visual targets or other random targets. Once the sequential organization “fixate-then-reach” has been learned (middle), reaching movements occur less frequently, because they are inhibited unless a fixation preceded reaching. Reaches tend to be successful, however. After the two behaviors fixate and reach have been boosted (right), reaches are more frequent and are directed at the visual targets.

5 Discussion

In this chapter, we have reviewed the five basic processes that are minimally required to direct movements at objects. To illustrate the style in which these processes may be neurally implemented, we have formulated and then simulated a simple neural dynamic model of movement generation. This model served as a platform for discussing the many developmental challenges that infants face as they learn to reach for objects.

Clearly, this model is as yet very rudimentary. Many components are mere place holders for neurally more realistic accounts. In many cases, more behavioral data may be needed to test or revise the assumptions we have made. We step through the

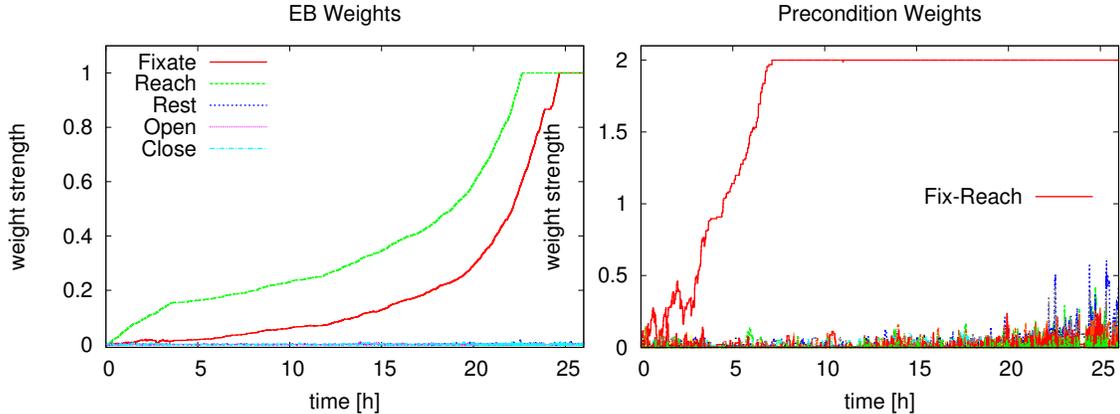


Figure 11: Evolution over learning time of the neural weights. Left: The boosts to elementary behaviors “fixate” and “reach” increase over time. Right: Only the boost to the pre-condition neural node for “fixate-then-reach” increases over learning time, all other pre-condition nodes remaining near zero.

components to discuss these various limitations and point to open questions and future directions of research.

5.1 Scene perception

We have only sketched the whole problem of scene perception and not modelled the underlying processes here (but see Zibner, Faubel, Iossifidis, and Schöner (2011)).

Looking behavior is extensively used to assess visual attention and perception in infants, both experimentally (Colombo & Mitchell, 2009) and theoretically (Perone & Spencer, 2013b; Schöner & Thelen, 2006). Typically, these studies make use of dwell time, that is, the duration of infants’ looks at particular parts of the visual array. In both habituation and preferential looking paradigms, the visual array is changed all the time. So if infants are building scene representations, those are constantly being challenged as the scene is changing. Scene perception in early infancy itself is not well studied and understood.

Scene representations must naturally be invariant against gaze shifts. A representation in retinal coordinates does not enable redirecting gaze or attention to an object from different viewing angles. Representing the visual array in a coordinate frame that is centered in the body enables such gaze shifts. The transformation from visual input in retinal coordinates to such a body-centered reference frame is neuronally costly (Pouget & Snyder, 2000; Schneegans, 2015). It requires projections from every retinal location to every possible location in a body reference frame. The current gaze angle selects the subset of these projections that reflects the current mapping at a given gaze direction. The demands on the neural machinery required to make coordinate transforms may be reflected in developmental processes that support the learning of such coordinate transforms (Sandamirskaya & Storck, 2015). At this time, little is known

about how infants learn coordinate transforms.

On the other hand, infants clearly need to know about the objects they reach to. The intriguing results of Corbetta et al. (2014) suggest, that while infants may be paying attention to an object they reach to on a global scale, the exact location on the object toward which they reach is not predicted by where they look. Is it possible, that learning to reach may impact on learning to look? Because the hand is attached to the body and the visual scene may be represented in a gaze-invariant, but body-related reference frame, it is thinkable that sensory and efference copy processes engaged in reaching contribute to the learning of spatial scene representations and of the associated coordinate transforms.

5.2 Movement preparation

We have assumed that reaching movements are prepared before they are initiated and that this preparation is based on information about both the movement target and the initial position of the hand. This is consistent with a broad range of empirical evidence, from neurophysiology to kinesiology and cognitive psychology: neural populations encoding the hand's movement direction in space, approximately straight movement paths of the hand in space with smooth and invariant velocity profiles, and the dependence of reaction times on movement parameters.

On the other hand, some data point toward a special role of the end-point of a movement compared to the initial posture. Stimulating small neural populations in motor cortex, Graziano and colleagues (Graziano, 2006) have often seen movements of the hand to particular locations on the body from wherever the hand happens to be initially, for instance. Infants learning to reach do not look at their hands prior to movement initiation and do not seem to pay attention to where their hands are (Corbetta et al., 2014). In fact, their hands seem to be all over the place as they attempt to reach repeatedly (Thelen et al., 1996). So is it possible that there would be an alternative approach to movement, one in which movement consists of moving directly to a new postural state irrespective of initial state and the spatial path required? We do not think that this is plausible. Timing and control problems are really quite different when the movement originates at other locations. A different angle may be this: Perhaps the current position of the hand is at all times neurally represented (as it is in our model). In Graziano's experiment, the activated neural populations first drive an update of a neural representation of the initial position of the hand based on this ongoing estimate. The same population of neurons may thus drive different movements depending on the initial position of the hand. This would be consistent with neural data from Scott and Kalaska (1997) who found that neural tuning curves in motor cortex depend on the arm's kinematic configuration.

Infants may have a difficulty organizing this sequence of processing steps. That difficulty may account for some part of the variance of their early reaches. Learning to reach may thus entail learning to temporally organize this sequence of processing steps (see Section 5.5).

5.3 Movement timing

Looking at infants' hand movements, it is obvious that one of their major difficulties is to get the timing of their reaches right. Some infants reach the target at high speeds, others make meek movements. Over development, these patterns of timing converge (Thelen et al., 1993, 1996).

We have talked about and modelled timing in an abstract way: A field of neural oscillators generates a virtual velocity profile that was assumed to be isomorphic in some way to the velocity of the hand's trajectory in space. Clearly, there is more to timing in movement generation. The time course of motor commands needs to be adapted and fine-tuned not only to assure coordination with other actions and perceived events, but also to support control, overcoming variations of inertia, interaction torques, or external forces acting on the moving limb. This is linked to the problem of optimal control: determining the right control signals that move an actuator to a desired state while satisfying constraints.

There is a huge literature on how adults learn to adapt motor commands so as to compensate for external force fields of a different kinds, including force fields that vary with velocity (Shadmehr & Mussa-Ivaldi, 2012). Interestingly, that experimental work has typically imposed force-fields on the end-effector, often the hand, rather than on individual joints (in part, for technical reasons, perhaps). An idea could be that an adaptation of the field of neural oscillators and of its projections onto the motor control level may be capable of modulating timing signals to accommodate the demands of these tasks. This is an open research question that may have impact for development because the changes in strength, mass, and geometry that occur as infants grow requires adaptation.

5.4 Motor control

Learning to time motor commands to reach targets is closely related to the issue of motor control, of course. In fact, researchers embracing the language of the theory of optimal control and of control theory do not typically break up the system into “timing” and “control” (although one could think of the optimal control signal as reflecting the timing level and the feedback controller as reflecting the control level).

One reason we propose to break up the system into “timing” and “control” is because of the particular way in which control of movement and control of posture are unified in human movement. The most general and consequential insight from the equilibrium point account of muscle neurophysiology is that due to the peripheral spinal circuitry (in particular, the stretch reflex) the descending control signal is essentially spatial in nature (Feldman, 2011). Muscles are activated as a function of how long they are (as sensed by muscle spindles, for instance) relative to a threshold. It is this threshold that is set by the descending motor command (and its rate of change), which can therefore be thought of as a length (and a velocity).

This solves a fundamental problem of muscle neurophysiology: After bursts of activity during a movement, muscles return to silence (or to an equilibrium level of

activity that may reflect stiffness demands and gravitational forces) even though the lengths of different muscles are not different than they were before the movement was started. So, the threshold of the stretch reflex must be reset during the movement. That requirement is sometimes left out of accounts that focus on the computation of the right torque profile that may overcome the various biomechanical interaction torques that plague movement. Computed torques as signals to the muscle are not only physiologically unrealistic, but also fail to address the problem of resetting the equilibrium length of muscles (which does not mean that motor commands may not be adjusted to deal with forces, e.g. (Tee, Burdet, Chew, & Milner, 2004)).

One hypothesis that we embrace, but have not elaborated here, is that the peripheral neural circuitry also simplifies the motor control problem (Gribble et al., 1998; Raphael, Tsianos, & Loeb, 2010). This may alleviate some of the problems of dealing with interaction torques or the torques experienced at individual joints due to external forces acting on the hand. A critical issue is how to resolve co-contraction and the redundancies among muscles systems.

The wager of equilibrium point thinking is that much of the remaining problem of control can be achieved through kinematic transformations that shape virtual joint trajectories (or muscle length trajectories). Kinematic transformations from an abstract “hand-in-space” virtual trajectory, $\mathbf{r}(t)$, to joint- or muscle level virtual trajectories, $\boldsymbol{\lambda}(t)$, may go a long way toward addressing control problems if we go beyond the simplistic ideas of inverse Jacobians that we employed here (Eq. 7). This is particularly attractive in a dynamical system (or neural dynamics) approach, in which the virtual joint trajectory is represented as the state of a dynamical system which may be used to express the configuration and velocity dependent modulation of motor commands, e.g.

$$\dot{\boldsymbol{\lambda}}(t) = \hat{\mathbf{J}}(\boldsymbol{\lambda}(t), \dot{\boldsymbol{\lambda}}(t)) \cdot \dot{\mathbf{r}}(t) \quad (16)$$

where the function, $\hat{\mathbf{J}}$, may serve as an internal model that must be learned. (Note that this implicit dynamical system needs to be resolved first.) Such a generalized kinematic transform could be thought of as a steerable neural mapping, analogous to the coordinate transforms we discussed earlier (Eq. 2). This mapping is from the spatial representation of the hand’s virtual velocity profile, $\dot{\mathbf{r}}(t)$ to the virtual joint level trajectory, $\boldsymbol{\lambda}(t)$ and $\dot{\boldsymbol{\lambda}}(t)$, steered by these same variables. Understanding how such kinematic transforms may be learned autonomously is one of the outstanding problems of developmental motor control (see Guenther, Bullock, Greve, and Grossberg (1994) for early attempts to achieve such).

This is linked to the “uncontrolled manifold” (UCM) effect, a signature of how the neural processes of movement generation deal with the inherent redundancy of the kinematics and dynamics of reaching (Scholz & Schöner, 1999; Tseng, Scholz, Schöner, & Hotchkiss, 2003). The UCM effect is the observation that there is more variance in those directions of joint space along which the hand position remains invariant than in directions along which the hand position varies. The UCM effect has been interpreted as revealing of a control strategy that focuses on the control of the hand in space (Todorov & Jordan, 2002; Latash et al., 2007; Martin et al., 2009). Such UCM structure of variance may thus serve to probe the structure of the kinematic

transformation. In particular, self-motion within the UCM is an index to incomplete compensation for interaction torques (Martin et al., 2009). Such analyses are difficult to perform for children, very difficult for infant reaching movement because many of the technical prerequisites of estimating the structure of variance are not met in these cases (e.g., temporal alignment of trajectories across different movement). It would be very attractive if these difficulties could be overcome.

5.5 Movement initiation and termination and learning to organize movements

The autonomous behavioral organization of reaching movement is a key problem, that is, perhaps, least well understood of all component processes. Behavioral organization is critical to integration, in which the component processes are initiated at the right time and hand on their output to downstream components when required. We have argued that sequentially organizing components of object-oriented action is a developmental achievement, referring to the data of von Hofsten (1984) for exemplary demonstration.

The critical issue is how the discrete moments in time, at which transitions among different components or different elementary movements occur, emerge from the underlying continuous time in which neural processes unfold. This is where the concepts of dynamical field theory provide a key element, the detection instability. At that instability, graded changes of input induced a categorical change that occurs at a discrete moment in time.

The emergence of such discrete events is also a key element of autonomous learning. We have argued that in autonomous learning, updates of neural connectivity only happens at particular points in time, at which a rewarding achievement is detected. Patterns of activation that lead up to such achievements must be kept around to properly assign credit at these learning events. We suggested that memory and eligibility traces were the simplest mechanism that may serve such function. Clearly, much more needs to be understood about autonomous learning.

We have also illustrated some of the processing structure that enables reaching movements, introducing the notions of intentional nodes, conditions of satisfaction, and of dissatisfaction. Learning requires even more prior structure, essentially the substrate to extract cues about the sequence of events, to link motor commands to their consequences, and to explore the behavioral options. Autonomous learning behavior from experience is, in our view, a major frontier of theoretical thinking in development (Perone & Spencer, 2013a; Sandamirskaya & Storck, 2015). We are only at the very beginning.

References

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

- Bastian, A., Schöner, G., & Riehle, A. (2003). Preshaping and continuous evolution of motor cortical representations during movement preparation. *European Journal of Neuroscience*, *18*, 2047–2058.
- Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *Journal of Neurophysiology*, *79*(6), 3168–3188.
- Berthier, N. E., & Keen, R. (2006). Development of reaching in infancy. *Experimental Brain Research*, *169*(4), 507–18.
- Botvinick, M., & Plaut, D. C. (2004). Doing Without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action. *Psychological Review*, *111*(2), 395–429.
- Brenner, E., Driesen, B., & Smeets, J. B. J. (2014). Precise timing when hitting falling balls. *Frontiers in human neuroscience*, *8*, 342.
- Bullock, D., Cisek, P., & Grossberg, S. (1998). Cortical Networks for Control of Voluntary Arm Movements under Variable Force Conditions. *Cerebral Cortex*, *8*, 48–62.
- Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychological Review*, *95*, 49–90.
- Caligiore, D., Parisi, D., & Baldassarre, G. (2014). Integrating reinforcement learning, equilibrium points, and minimum variance to understand the development of reaching: A computational model. *Psychological Review*, *121*(3), 389–421.
- Chao, F., Lee, M., & Lee, J. (2010). A developmental algorithm for ocularmotor coordination. *Robotics and Autonomous Systems*, *58*(3), 239–248.
- Chevalier, G., & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in Neurosciences*, *13*(7), 277–280.
- 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., Grossberg, S., & Bullock, D. (1998). A cortico-spinal model of reaching and proprioception under multiple task constraints. *Journal of Cognitive Neuroscience*, *10*(4), 425–444.
- 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.
- Colombo, J., & Mitchell, D. (2009). Infant visual habituation. *Neurobiology of learning and memory*, *92*(2), 225–234.
- Corbetta, D., & Thelen, E. (1995). A method for identifying the initiation of reaching movements in natural prehension. *Journal of Motor Behavior*, *27*, 285–293.
- Corbetta, D., Thurman, S. L., Wiener, R. F., Guan, Y., & Williams, J. L. (2014). Mapping the feel of the arm with the sight of the object: on the embodied origins of infant reaching. *Frontiers in psychology*, *5*, 576.
- Erlhagen, W., & Schöner, G. (2002). Dynamic field theory of movement preparation. *Psychological Review*, *109*(3), 545–572.

- Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics*, *61*, 353–430.
- Fard, F. S., Hollensen, P., Heinke, D., & Trappenberg, T. P. (2015). Modeling human target reaching with an adaptive observer implemented with dynamic neural fields. *Neural networks : the official journal of the International Neural Network Society*, *72*, 13–30.
- Feldman, A. G. (1986). Once more on the equilibrium point hypothesis (λ -model) for motor control. *Journal of Motor Behavior*, *18*, 15–54.
- Feldman, A. G. (2011, may). Space and time in the context of equilibrium-point theory. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(3), 287–304.
- Gallistel, C. R., & Gibbon, J. (2000). Time, rate, and conditioning. *Psychological Review*, *107*(2), 289–344.
- Georgopoulos, A. P. (1986). On reaching. *Annual Reviews of Neuroscience*, *9*, 147–170.
- Ghafouri, M., & Feldman, A. G. (2001). The timing of control signals underlying fast point-to-point arm movements. *Experimental Brain Research*, *137*, 411–423.
- Graziano, M. S. A. (2006). The organization of behavioral repertoire in motor cortex. *Annual Reviews of Neuroscience*, *29*, 105–134.
- Gribble, P. L., Ostry, D. J., Sanguineti, V., & Laboissière, R. (1998). Are complex control signals required for human arm movement. *Journal of Neurophysiology*, *79*, 1409–1424.
- Grimme, B., Lipinski, J., & Schöner, G. (2012). Naturalistic arm movements during obstacle avoidance in 3D and the identification of movement primitives. *Experimental brain research*, *222*(3), 185–200.
- Guenther, F., Bullock, D., Greve, D., & Grossberg, S. (1994). Neural representations for sensory-motor control. III. Learning a body-centered representation of 3-D target position. *Journal of Cognitive Neuroscience*, *6*, 341–358.
- Herbort, O., & Butz, M. (2009). The SURE-REACH Model for Motor Learning and Control of a Redundant Arm: From Modeling Human Behavior to Applications in Robotics. *From Motor Learning to Interaction Learning in Robots*, 85–106.
- Hermer, L., & Spelke, E. S. (1994). A geometric process for spatial orientation in young children. *Nature*, *370*, 57–59.
- Hollingworth, A. (2004). Constructing visual representations of natural scenes: The roles of short- and long-term visual memory. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 519–537.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*(10-12), 1489–1506.
- Jeannerod, M. (1984). The timing of natural prehension movements. *Journal of Motor Behavior*, *16*, 235–254.
- Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A Dynamic Neural Field Model of Visual Working Memory and Change Detection. *Psychological Science*, *20*, 568–577.
- Kazerounian, S., Luciw, M., Richter, M., & Sandamirskaya, Y. (2013). Autonomous Reinforcement of Behavioral Sequences in Neural Dynamics. In *Proceedings of*

- the international joint conference on neural networks (IJCNN 2013).*
- Kelso, J. A. S., Southard, D. L., & Goodman, D. (1979). On the nature of human interlimb coordination. *Science*, *203*, 1029–1031.
- Kopecz, K. (1995). Saccadic reaction times in gap/overlap paradigms: a model based on integration of intentional and visual information on neural, dynamic fields. *Vision Research*, *35*, 2911–2925.
- Kopecz, K., & Schöner, G. (1995). Saccadic motor planning by integrating visual information and pre-information on neural dynamic fields. *Biological Cybernetics*, *73*, 49–60.
- Latash, M. L. (2008). *Synergy*. New York: Oxford University Press.
- Latash, M. L., Scholz, J. P., & Schöner, G. (2007). Toward a New Theory of Motor Synergies. *Motor Control*, *11*, 276–308.
- Lisman, J. (2015). The Challenge of Understanding the Brain: Where We Stand in 2015. *Neuron*, *86*(4), 864–882.
- Loram, I. D., & Lakie, M. (2002). Direct measurement of human ankle stiffness during quiet standing: the intrinsic mechanical stiffness is insufficient for stability. *Journal of Physiology*, *545*, 1041–1053.
- Martin, V., Scholz, J. P., & Schöner, G. (2009). Redundancy, self-motion and motor control. *Neural Computation*, *21*(5), 1371–1414.
- McDowell, K., Jeka, J. J., Schöner, G., & Hatfield, B. D. (2002). Behavioral and electrocortical evidence of an interaction between probability and task metrics in movement preparation. *Experimental brain research*, *144*(3), 303–13.
- Mechsner, F., Kerzel, D., Knoblich, G., & Prinz, W. (2001). Perceptual basis of bimanual coordination. *Nature*, *414*, 69–73.
- Moran, D. W., & Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching movement. *Journal of Neurophysiology*, *82*, 2676–2692.
- Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, *42*, 223–227.
- Narioka, K., & Steil, J. J. (2015). U-shaped motor development emerges from Goal Babbling with intrinsic motor noise. *Proceedings of the 5th Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics ICDL-EpiRob 2015*, 55–62.
- Perone, S., & Spencer, J. P. (2013a). Autonomous visual exploration creates developmental change in familiarity and novelty seeking behaviors. *Frontiers in Psychology*, *4*(648), 1–21.
- Perone, S., & Spencer, J. P. (2013b). Autonomy in action: linking the act of looking to memory formation in infancy via dynamic neural fields. *Cognitive science*, *37*(1), 1–60.
- Pezzulo, G., & Cisek, P. (2016). Navigating the Affordance Landscape: Feedback Control as a Process Model of Behavior and Cognition. *Trends in Cognitive Sciences*, *20*(6), 414–424.
- Pouget, A., & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature*, *3*, 1192–1198.
- Raphael, G., Tsianos, G. a., & Loeb, G. E. (2010). Spinal-like regulator facilitates

- control of a two-degree-of-freedom wrist. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 30(28), 9431–44.
- Richter, M., Sandamirskaya, Y., & Schöner, G. (2012). A robotic action selection and behavioral organization architecture inspired by human cognition. In *Ieee/rsj international conference on intelligent robots and systems* (pp. 2457–2464).
- Ruf, A., & Horaud, R. (1999). Visual Servoing of Robot Manipulators Part I: Projective Kinematics. *The International Journal of Robotics Research*, 18(11), 1101–1118.
- Sandamirskaya, Y. (2015). Autonomous Sequence Generation in Dynamic Field Theory. In G. Schöner, J. P. Spencer, & T. DFT Research Group (Eds.), *Dynamic thinking: A primer on dynamic field theory* (pp. 353–368). New York, NY, USA: Oxford University Press.
- Sandamirskaya, Y., & Conradt, J. (2013). Learning transformations with Dynamic Neural Fields. In *Icann 2013*.
- Sandamirskaya, Y., & Schöner, G. (2010). An embodied account of serial order: how instabilities drive sequence generation. *Neural networks*, 23(10), 1164–79.
- Sandamirskaya, Y., & Storck, T. (2015). Learning to Look and Looking to Remember: A Neural-Dynamic Embodied Model for Generation of Saccadic Gaze Shifts and Memory Formation. In P. Koprinkova-Hristova, V. Mladenov, & N. K. Kasabov (Eds.), *Artificial neural networks se - 9* (Vol. 4, pp. 175–200). Springer International Publishing.
- Savastano, P., & Nolfi, S. (2013). A robotic model of reaching and grasping development. *IEEE Transactions on Autonomous Mental Development*, 5(4), 326–336.
- Schlesinger, M., Parisi, D., & Langer, J. (2000). Learning to Reach by Constraining the Movement Search Space. *Developmental Science*, 3, 67–80.
- Schneegans, S. (2015). Sensori-Motor and Cognitive Transformation. In G. Schöner, J. P. Spencer, & T. DFT Research Group (Eds.), *Dynamic thinking: A primer on dynamic field theory* (pp. 169–196). New York, NY, USA: Oxford University Press.
- Scholz, J. P., & Schöner, G. (1999). The Uncontrolled Manifold Concept: Identifying Control Variables for a Functional Task. *Experimental Brain Research*, 126, 289–306.
- Schöne, H. (1984). *Spatial Orientation*. The spatial control of behavior in animals and man. Princeton, NJ: Princeton University Press.
- Schöner, G. (1990). A dynamic theory of coordination of discrete movement. *Biological Cybernetics*, 63, 257–270.
- Schöner, G. (2002). Timing , Clocks , and Dynamical Systems. *Brain and Cognition*, 48, 31–51.
- Schöner, G., Spencer, J. P., & the DFT Research Group. (2015). *Dynamic Thinking: A Primer on Dynamic Field Theory*. Oxford University Press.
- Schöner, G., & Thelen, E. (2006). Using Dynamic Field Theory to Rethink Infant Habituation. *Psychological Review*, 113(2), 273–299.
- Schwartz, A. B. (1994). Direct Cortical Representation of Drawing. *Science*, 265, 540–542.
- Scott, S. H., & Kalaska, J. F. (1997). Reaching movements with similar hand paths but

- different arm orientations. I. Activity of individual cells in motor cortex. *Journal of Neurophysiology*, *77*, 826–852.
- Shadmehr, R., & Krakauer, J. W. (2008). A computational neuroanatomy for motor control. *Experimental Brain Research*, *185*(359-381).
- Shadmehr, R., & Mussa-Ivaldi, F. (2012). *Biological Learning and Control: how the brain builds representations, predicts events, and makes decisions*. Cambridge, MA, USA: MIT Press.
- Simons, D. J. (2000, jan). Current Approaches to Change Blindness. *Visual Cognition*, *7*(1-3), 1–15.
- Soechting, J. F., & Lacquaniti, F. (1981). Invariant characteristics of a point movement in man. *Journal of Neuroscience*, *1*(7), 710–720.
- Sporns, O., & Edelman, G. M. (1993). Solving Bernstein’s problem: a proposal for the development of coordinated movement by selection. *Child development*, *64*(4), 960–981.
- Strauss, S., & Heinke, D. (2012). A Robotics-Based Approach to Modeling of Choice Reaching Experiments on Visual Attention. *Frontiers in Psychology*, *3*, 1–11.
- Strauss, S., Woodgate, P. J., Sami, S. A., & Heinke, D. (2015). Choice reaching with a LEGO arm robot (CoRLEGO): The motor system guides visual attention to movement-relevant information. *Neural Networks*, *72*, 3–12.
- Sun, G., & Scassellati, B. (2005). A fast and efficient model for learning to reach. *International Journal of Humanoid Robotics*, *2*(4), 391–414.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundation of adaptive networks* (pp. 497–537). MIT Press, Cambridge MA, USA.
- Tee, K. P., Burdet, E., Chew, C. M., & Milner, T. E. (2004). A model of force and impedance in human arm movements. *Biological cybernetics*, *90*, 368–375.
- Tekülve, J., Zibner, S. K. U., & Schöner, G. (2016). A neural process model of learning to sequentially organize and activate pre-reaches. *Proceedings of the 6th Joint IEEE International Conferences on Development and Learning and on Epigenetic Robotics (ICDL-Epirob) 2016*.
- Thelen, E. (1981). Rhythmical behavior in infancy: An ethological perspective. *Developmental Psychology*, *17*(3), 237–257.
- Thelen, E., Corbetta, D., Kamm, K., Spencer, J., Schneider, K., & Zernicke, R. (1993). The transition to reaching: Matching intention and intrinsic dynamics. *Child Development*, *64*, 1058–1098.
- Thelen, E., Corbetta, D., & Spencer, J. P. (1996). The development of reaching during the first year: The role of movement speed. *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 1059–1076.
- Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–1236.
- Tseng, Y.-W., Scholz, J. P., Schöner, G., & Hotchkiss, L. (2003). Effect of accuracy constraint on joint coordination during pointing movements. *Experimental Brain Research*, *149*(3), 276–288.

- Viviani, P., & Flash, T. (1995). Minimum-jerk, two-thirds power law, and isochrony: converging approaches to movement planning. *Journal of experimental psychology. Human perception and performance*, *21*(1), 32–53.
- von Hofsten, C. (1984). Developmental changes in the organization of prereaching movements. *Developmental Psychology*, *20*(3), 378–388.
- von Hofsten, C. (1991). Structuring of early reaching movements: a longitudinal study. *Journal of motor behavior*, *23*(4), 280–292.
- Wilimzig, C., Schneider, S., & Schöner, G. (2006). The time course of saccadic decision making : Dynamic field theory. *Neural Networks*, *19*, 1059–1074.
- Zibner, S. K. U., Faubel, C., Iossifidis, I., & Schöner, G. (2011). Dynamic Neural Fields as Building Blocks for a Cortex-Inspired Architecture of Robotic Scene Representation. *IEEE Transactions on Autonomous Mental Development*, *3*(1), 74–91.
- Zibner, S. K. U., Tekülve, J., & Schöner, G. (2015a). The neural dynamics of goal-directed arm movements : a developmental perspective. In *Proceedings of the 5th joint international conference on learning and development and on epigenetic robotics ICDL-EpiRob 2015*.
- Zibner, S. K. U., Tekülve, J., & Schöner, G. (2015b). The Sequential Organization of Movement is Critical to the Development of Reaching : A Neural Dynamics Account. In *Proceedings of the 5th joint international conference on learning and development and on epigenetic robotics ICDL-EpiRob 2015*.