The neural dynamics of goal-directed arm movements: a developmental perspective

Stephan K. U. Zibner, Jan Tekülve, and Gregor Schöner
Institut für Neuroinformatik, Ruhr-Universität Bochum
Universitätsstr. 150, 44780 Bochum, Germany
Email: {stephan.zibner, jan.tekuelve, gregor.schoener}@ini.rub.de

***Abstract*—We present a neuro-dynamic architecture for the generation of movement of the hand toward a visual target that integrates movement planning based on visual input, movement initiation and termination, the generation of the time courses of virtual trajectories of the hand in Cartesian space, and their transformation into virtual joint trajectories and muscle forces. The architecture captures properties of adult goal-directed arm movements such as bell-shaped velocity profiles and on-line updating of a movement when the target is shifted. The integrated and autonomous nature of the architecture makes it possible to study how motor performance is affected when one of the three core processes, planning, timing, and transformation from end-effector to joint space, are decalibrated to reflect earlier stages of development. We find signatures of the development of reaching such as multiple movement units and curved movement paths in the "young" model.**

***Index Terms*—movement generation, onset of reaching, dynamic field theory, equilibrium point theory.**

I. Introduction

Understanding how the human nervous system generates goal-directed arm movements continues to be a challenging research topic. A central difficulty is that movement generation is a highly integrated process, in which perception (e.g., of the movement target or of obstacles), cognition (e.g., planning of movement, remembering movement targets), and movement generation (e.g., timing movement and activating muscular processes) are tightly coupled. One signature of such tight integration is the phenomenon of on-line updating, in which a movement aimed at a visual target can be adjusted any time during movement preparation or execution with a delay of about 100 ms to a change in target location [33], [8], [7]. On-line updating does not require awareness of the target change, which is tested by shifting the target during a saccadic eye movement [17]. Thus, movement plans are continuously linked to visual input and continuously steer the movement of the hand.

The integrated nature of movement generation is not just a problem for us, the scientist, trying to decompose the systems. It is also a problem for the nervous system as it learns to generated goal-oriented movements. Learning to reach is a critical developmental milestone and infants work very hard at it. Many processes need to be in place in order to successfully reach: The mapping from visual space to a body-centered motor space must work [5]. Body-centered information about the target or movement plan must be translated into joint and muscle space. Time courses of motor commands must be generated that bring about movement of the arm to the target, which must be reached with low enough speed so that soft contact is possible [29]. Clearly, these processes strongly interrelate. Autonomously learning to make targeted movements requires that these processes be aligned [20]. For instance, to know if the gain of a motor map is too large or too small, visual information about the hand in space must be related back to the representation of the motor command. How reaching develops is hard to study for the same reason: many components may change at the same time. Two salient and robust developmental signatures of the onset of reaching behavior are, however, the increasing straightness of hand paths in space [29] and the reduction of the number of movement units (MUs, number of significant maxima in hand speed during the movement [34]).

In this paper we study an architecture that generates goal-directed arm movements and integrates the entire neural processing pathway from sensory input to motor output. We use dynamic field theory (DFT) [23], a theoretical framework for neural dynamics, throughout the architecture. The model is based on three commitments. First, we consider movement planning to take place at the level of the end-effector's (the hand's) movement in space. This is consistent with behavioral signatures such as the bell-shaped velocity profiles of the hand in space, or the invariant shape of hand paths at different movement speeds and loads [16], [1]. The structure of joint variance is another source of evidence for the special role of the spatial path of the hand in movement planning. Joint configurations that leave the hand position in space invariant fluctuate more from trial to trial than joint configurations that change the hand position in space [32]. This assumption is also consistent with neural evidence in which neural populations in motor and pre-motor cortex represent the direction of the hand's movement in space [9] and its amplitude [15].

Our second commitment is to consider the timing of goal-directed arm movements to originate from neural oscillators that generate virtual hand trajectories. Coupled oscillators have emerged as a theoretical framework to understand much behavioral evidence about the coordination of rhythmic [24] and discrete [21] movement (reviewed in [12], [22]). That these oscillators correlate with the spatial representation of movement trajectories is consistent with much of that evidence, including targeted experiments that realize the same

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

spatial pattern of coordinated movements with different effector combinations [14]. This commitment is also consistent with the discovery of an oscillatory component of neural population activity in motor cortex [4].

Our third commitment is to take into account the dynamics of force generation by muscles. Spinal feedback loops couple muscle force generation to muscle length. As a result, descending motor commands, λ, are essentially spatial in nature: they set the equilibrium muscle length in the presence of a given level of external force [6]. Forces are generated in response to the deviation of the current muscle length from the equilibrium length. The processes of torque generation is characterized by delays induced by its own dynamics [10], which further complicate the task of generating the torques required to move an effector to a target.

What sets our approach apart from related modeling work is the autonomy of all processes required to generate targeted movements. Through a system of neural dynamics, the architecture autonomously initiates movement when the visual target is perceived far from where the hand rests. The system terminates motion autonomously when that is no longer the case. We share this detection of a distance between current and target state with the classical neural dynamic model of Bullock and Grossberg [3]. In that model, however, movement is triggered by a go signal that comes from the outside. This makes it difficult to understand how multiple motor units may emerge autonomously. Optimal control models are also autonomous only to a limited extent in that they compute the optimal motor command but do not address how movements are started or stopped and how the updating of the movement occurs as a process when the target shifts [31]. The model closest to ours [28] shares the commitment to a spatial representation of motor plans and of movement timing, and is technically in the same language of dynamic field theory. It differs in the details of the timing mechanism and does not address the muscular level. Moreover, its autonomy is also less comprehensive. Rokni and Sompolinsky [19] present a model that overlaps for the second two commitments. It uses a bank of neural oscillators with tuneable frequencies and amplitudes to generate movement timing and uses a simple muscle model to account for the control of the plant. This model does not explicitly include the planning level and thus is not open to understanding on-line updating. It is also limited in its autonomy.

We implement the architecture for an anthropomorphic robot arm with seven degrees of freedom that reaches to targets localized in a two-dimensional plane. We target the CAREN platform, a stationary robot consisting of a camera mounted on a two degrees-of-freedom head and a seven degrees-of-freedom Kuka light-weight robot (LWR 4) [13]. In this paper, the architecture is evaluated in simulation. We emphasize on-line updating as a signature of the integrated nature of the model. We exploit the capacity of the model to autonomously organize movement even when the hand does not reach the target to study the role of three key component processes. According to the spatial precision hypothesis, earlier in development, neural interaction is weaker leading to smaller capacity of working memory [30], [26], [25], [27]. We look at how this "younger" dynamic regime in the movement planning system affects movement. We also examine how changes to the gain of the movement timing system as well as to the inverse kinematics affect motor performance and compare the resultant movement patterns to developmental findings [29].

II. Methods

In dynamic field theory, dynamic neural fields (DNFs) and dynamic neural nodes are the building blocks of neural architectures. The dynamics of a neural activation field, $u(\boldsymbol{x}, t)$, defined over d dimensions, $\boldsymbol{x} = (x_1, \ldots, x_d)^{\mathrm{T}}$:

$$\begin{aligned} \tau \dot{u}(\boldsymbol{x}, t) &= -u(\boldsymbol{x}, t) + h + [w_{u,u} * \sigma(u)](\boldsymbol{x}, t) \\ &\quad + \sum_i s_i(\boldsymbol{x}, t) \end{aligned} \tag{1}$$

has attractor states that are shaped by the homogeneous resting level, h, by local excitatory and global inhibitory interactions described by the bell-shaped interaction kernel $w_{u,u}(\boldsymbol{x})$, and by the sum of all inputs, s_i, to the field. Interaction is mediated by a nonlinear sigmoidal output function,

$$\sigma(u(\boldsymbol{x}, t)) = 0.5 \left(1 + \frac{\beta u(\boldsymbol{x}, t)}{1 + \beta |u(\boldsymbol{x}, t)|}\right), \tag{2}$$

which is convolved with the interaction kernel as indicated by the symbol $*$. Inputs may originate in other layers of our architecture or in the sensory surface itself (e.g., the camera image). For a more detailed analysis of the dynamics of DNFs and the possible types of couplings between different fields, see [23], [35].

Peaks of supra-threshold activation represent estimates of the state vector, $\boldsymbol{x}$, determined by the peak's position along the field's dimensions. Such peaks may arise as attractor states from instabilities of the sub-threshold patterns of activation, for instance, in response to localized input (detection instability). Peaks may become unstable when localized input is removed or when inhibition impacts that may originate from competing peaks or from inhibitory inputs. DNFs may thus make selection decisions. When lateral excitation is sufficiently strong, peaks may become sustained so that they remain stable when inducing localized input is removed. Match detection can be implemented by projecting two DNFs onto a third DNF, which generates a peak where inputs overlap. Mismatch detection can be implemented in the same way by making one of the projections inhibitory [11].

The second type of building block of DFT architectures are dynamic neural nodes (that can be thought of as DNFs with homogeneous activation patterns):

$$\tau \dot{u}(t) = -u(t) + h + c_{u,u} \sigma(u(t)) + \sum_i s_i(t). \tag{3}$$

The interaction kernel is replaced by a scalar connection weight, $c_{u,u}$, the strength of self-excitation. The output function, $\sigma(\cdot)$, is the same as defined previously in Equation 2.

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

A node is considered *active* when its activation is above the threshold of the sigmoid function. This state arises as a fixed point attractor from a detection instability analogously to that of DNFs. In DFT architectures, nodes are used to represent intentional states and their conditions of satisfaction [18]. Intention nodes project onto DNFs as global boosts of field activation that drive fields through the detection instability, leading to the creation of a peak. Intention nodes may also project locally (through a forward kernel) to provide a competitive advantage to particular regions in a field. In either case, intention nodes promote the creation of peaks that then specificy the intended behavior (e.g., a particular movement direction in a motor plan). Condition of satisfaction (CoS) nodes receive input both from an intention node and from a CoS field. Given input from the intention node, a peak in CoS field drives the CoS node above threshold. CoS nodes thus act as *peak detectors*. CoS fields detect the match between the intended behavior (represented in a DNF) and sensory input that signals the outcome of the intended action (represented in another DNF). A condition of dissatisfaction (CoD) can be defined analogously through mis-match detection. It marks the failure of a behavior and induces error recovery behaviors or a restart of the current behavior. These are issues of behavioral organization discussed in more detail elsewhere [18].

III. Architecture

We give a short overview of the DFT architecture for movement generation and then step through the components. A movement target is visually perceived and represented in a DNF (Figure 1, A). Using the initial state of the end-effector represented in a DNF (Figure 1, B), a movement plan is generated by transforming the target representation into coordinates centered on the initial position of the end-effector. This movement plan activates a particular neural oscillator within a two-layer field of neural activation (Figure 1, C). The oscillator generates a virtual end-effector velocity vector, $\boldsymbol{v}(t)$, that is transformed to a virtual joint velocity vector, $\dot{\boldsymbol{\lambda}}(t)$, of the seven degrees-of-freedom robot arm by applying an inverse kinematics. An integrator (Figure 1, D) passes the virtual joint trajectory, $\boldsymbol{\lambda}(t)$, on to the muscle model (Figure 1, E), which in turn drives the robot arm if the current state of the actuator differs from the equilibrium configuration of all muscles. In a form of corollary discharge, a second pathway generates a virtual end-effector trajectory from the virtual joint trajectory by forward kinematics, and transforms this into input to a DNF that represents the current state of the virtual end-effector (Figure 1, F). That field updates the neural representation of the initial position of the end-effector (Figure 1, B) whenever a CoS signals the end of a movement unit (CoS and CoD projections show in red in Figure 1). The entire architecture is a continuous time dynamical system, from which the discrete events that activate and deactivate components emerge autonomously.

A. Generation of a movement plan

The pathway from the visual sensor (an RGB camera) to a representation of the movement target is done in the style of [35]. A saliency operation transforms the image into an activation field, highlighting potential target candidates selected by a color cue. A reference frame transformation expresses the perceived retinal locations of objects in a reference frame that is linked to the two-dimensional coordinate system of the table in front of the robot. Note that any allocentric or body-centered reference frame is suitable as long as both the target and the end-effector can be expressed in the same reference frame and a transformation to motor commands is known. Only a subregion of the table plane is represented in the following DNFs (see Figure 1, top). Movement can only take place to targets inside this region. The target object is represented in a single-peak two-dimensional DNF, u_{tar}, (see Figure 1, A), following Equation 1, its dimensions matching the two-dimensional reference frame, $\boldsymbol{x} = (x_1, x_2)^{\mathrm{T}}$, of the table surface. Its sole input, $s_{\mathrm{tar}}(\boldsymbol{x}, t)$, originates in the bottom-up processing of the camera input. u_{tar} contains a peak if a region with strong activation exists in s_{tar}. The DNF's lateral interaction kernel assures through global inhibition that there is at most one peak at any time, preventing the architecture from pursuing multiple movement goals at the same time. This DNF tracks changes in the target's position.

The end-effector position is represented in two DNFs. One is a working memory of the end-effector position at the beginning of each movement, u_{ini}, a two-dimensional, single-peak DNF (Figure 1, B) The other is a neural representation, u_{cur}, of the current prediction of the end-effector position obtained from the outflowing motor command, $\boldsymbol{\lambda}$, using forward kinematics. The current prediction is used to update the initiation position whenever the virtual end-effector is resting by providing input to u_{ini}:

$$s_{\mathrm{ini}}(\boldsymbol{x}, t) = [w_{\mathrm{ini,cur}} * \sigma(u_{\mathrm{cur}})](\boldsymbol{x}, t). \tag{4}$$

Both fields are expressed in a two-dimensional reference frame spanning the same table surface over which u_{tar} is defined. The lateral interaction kernel and resting level of u_{ini} are set up to keep a single sustained activation peak as long as there is no significant level of activation in u_{cur}. u_{cur} is inhibited during movement, so that u_{ini} is updated only after a movement has finished.

A movement plan is calculated by transforming the target representation to a coordinate frame centered in the end-effector. The transformation is performed by convolving the output of the target field with the output of u_{ini},

$$s_{\mathrm{pla}}(\boldsymbol{x}, t) \quad = \quad [\sigma(u_{\mathrm{tar}}) * \sigma(u'_{\mathrm{ini}})](\boldsymbol{x}, t), \tag{5}$$

where the prime indicates that the dimensions are inverted. Such a transformation is a neural implementation of the subtraction of the end-effector position from the object position proposed by Bullock and Grossberg [3].

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

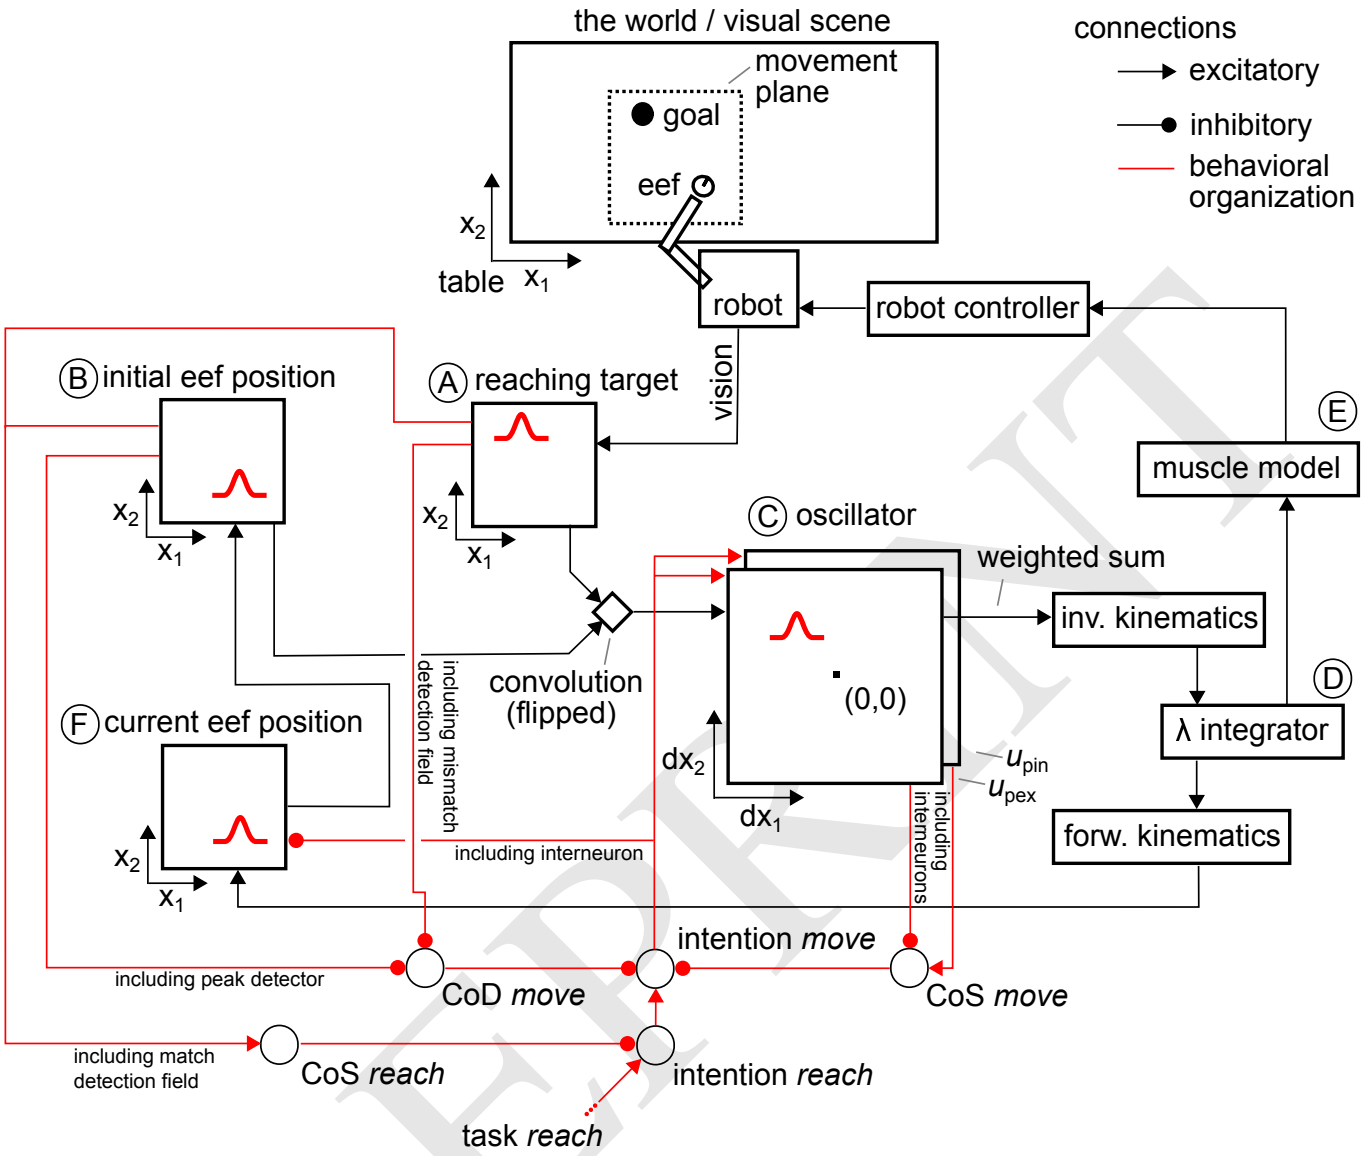

Fig. 1. This figure shows the full movement generation architecture. Some details are hidden in connections for clarity's sake, but are marked with text stating "including …". See text for more details.

B. Generation of virtual trajectory

The movement plan feeds into a two-layer DNF, consisting of u_{pex} and u_{pin} (see Figure 1, C),

$$\begin{aligned}\tau_{\text{pex}}\dot{u}_{\text{pex}}(\boldsymbol{x},t) &= -u_{\text{pex}}(\boldsymbol{x},t) + h + s_{\text{pex}}(\boldsymbol{x},t) \\ &\quad -[w_{\text{pex,pin}} * \rho(u_{\text{pin}})](\boldsymbol{x},t) \qquad (6)\\ \tau_{\text{pin}}\dot{u}_{\text{pin}}(\boldsymbol{x},t) &= -u_{\text{pin}}(\boldsymbol{x},t) + h + s_{\text{pin}}(\boldsymbol{x},t), \qquad (7)\end{aligned}$$

with $s_{\text{pex}}(\boldsymbol{x},t) = s_{\text{pin}}(\boldsymbol{x},t) = s_{\text{pla}}(\boldsymbol{x},t) + c_{\text{mov}}\sigma(u^{\text{int}}_{\text{mov}}(t))$ and $\tau_{\text{pex}} < \tau_{\text{pin}}$. The two-layer structure of u_{pex} and u_{pin} serves as a neural oscillator. Transient activation is created in the excitatory layer, which the more slowly evolving inhibitory layer suppresses over time. This dynamics thus performs a one-shot active transient in response to input. The oscillation is parameterized by the movement plan s_{pla} and is switched on by the activation of a dynamic neural node $u^{\text{int}}_{\text{mov}}$, which expresses the intention to generate movement. Both layers use a semi-linear output function $\rho(\cdot)$ instead of $\sigma(\cdot)$,

$$\rho(\boldsymbol{x},t) = \begin{cases} u(\boldsymbol{x},t) & \text{for } u(\boldsymbol{x},t) > 0 \\ 0 & \text{else.} \end{cases} \qquad (8)$$

This assures that no movement is created as long as u_{pex} is below threshold. Note that u_{pex} and u_{pin} cover a larger spatial area than u_{tar} and u_{ini}, as their coordinate system expresses relative distance to the end-effector. Consequently, if the end-effector is at the target, the target appears in the center of u_{pex} and u_{pin} with a distance of zero to the end-effector.

From the relative position of the target in u_{pex}, a velocity vector $\boldsymbol{v}$ is extracted by integrating over the represented domain $X = \{(x_1, x_2) \in \mathbb{R}^2 : -50 \le x_1, x_2 \le 50\}$:

$$\boldsymbol{v}(t) = \iint_X \rho(u_{\text{pex}}(\boldsymbol{x},t))\omega(\boldsymbol{x})\, dx_1 dx_2. \qquad (9)$$

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

$\omega(\cdot)$ is a weight function learned with gradient descent. These weights are a linear function of x_1 and x_2. They can be interpreted as tuning curves of the oscillator field's sites, with each site being sensitive to both the orientation and the distance (and thus peak velocity) of the movement.

C. Generation of muscle activation and arm movement

The virtual velocity vector, $\boldsymbol{v}$, is transformed into a virtual joint velocity vector, $\dot{\boldsymbol{\lambda}}$, using the pseudo-inverse of the manipulator Jacobian, $\boldsymbol{J}$, that depends on the current joint configuration $\boldsymbol{\theta}(t)$:

$$\dot{\boldsymbol{\lambda}} = \boldsymbol{J}^{+}(\boldsymbol{\theta}(t))\boldsymbol{v}(t), \tag{10}$$

from which the virtual joint configuration, $\boldsymbol{\lambda}(t)$, is obtained by integrating over time (Figure 1, D). $\boldsymbol{\lambda}(t)$ is subsequently passed to a linear second order muscle model, which combines agonist and antagonist muscles around each joint (Figure 1, E),

$$\ddot{\boldsymbol{\theta}} = -\boldsymbol{K}(\boldsymbol{\theta} - \boldsymbol{\lambda}) - \boldsymbol{B}\dot{\boldsymbol{\theta}}, \tag{11}$$

with $\boldsymbol{\theta}$ being the seven-dimensional joint configuration of the robot arm. $\boldsymbol{K}$ and $\boldsymbol{B}$ are matrices determining the stiffness and viscosity of the muscles. $\dot{\boldsymbol{\theta}}$ is sent to the robot controller using velocity control.

D. Corollary discharge

The virtual joint configuration, $\boldsymbol{\lambda}$, is transformed into input for the DNF u_{cur} (Figure 1, F):

$$s_{\mathrm{cur}}(\boldsymbol{x}, t) = s_{\mathrm{fwk}}(\boldsymbol{x}, t) - c_{\mathrm{cur,mov}}\sigma(u^{\mathrm{int}}_{\mathrm{mov}}(t)). \tag{12}$$

s_{fwk} is created by applying the forward kinematics to $\boldsymbol{\lambda}$, obtaining a predicted end-effector position, $\boldsymbol{x}^* = (x_1^*, x_2^*)^{\mathrm{T}}$,

$$s_{\mathrm{fwk}}(\boldsymbol{x}, t) = c \exp\left(-\left(\frac{(x_1 - x_1^*)^2}{2\sigma_{x_1}^2} + \frac{(x_2 - x_2^*)^2}{2\sigma_{x_2}^2}\right)\right) \tag{13}$$

where, $\boldsymbol{x} = (x_1, x_2)^{\mathrm{T}}$. The contribution $\sigma(u^{\mathrm{int}}_{\mathrm{mov}}(t))$ is used to de-boost u_{cur}, which puts u_{ini} from input-driven to self-sustained mode. u_{cur} can be interpreted as a forward model of the end-effector position so that a peak in u_{cur} represents the location to which the end-effector will move over time, delayed by the muscle model.

E. Behavioral organization

Behavioral organization takes care of initiating the generation of a virtual trajectory, $\boldsymbol{v}$, and of updating the initial end-effector position represented in u_{ini}. It does so by boosting u_{pex} and u_{pin} and de-boosting u_{cur} at the start of each virtual movement and taking away these boosts at the end of the virtual movement. An intention node, $u^{\mathrm{int}}_{\mathrm{mov}}$, is connected to these three fields. It receives input from three nodes,

$$\begin{aligned} s^{\mathrm{int}}_{\mathrm{mov}}(t) \;=\; & c_{\mathrm{mov,rea}}\sigma(u^{\mathrm{int}}_{\mathrm{rea}}(t)) \\ & - c_{\mathrm{int,CoS}}\sigma(u^{\mathrm{CoS}}_{\mathrm{mov}}(t)) - c_{\mathrm{int,CoD}}\sigma(u^{\mathrm{CoD}}_{\mathrm{mov}}(t)). \end{aligned} \tag{14}$$

It is activated by an intention node on a higher hierarchical level ($u^{\mathrm{int}}_{\mathrm{rea}}(t)$, see below) and inhibited by either a CoS or a CoD node. A CoS node $u^{\mathrm{CoS}}_{\mathrm{mov}}$ monitors the field activity in both u_{pex} and u_{pin}, signaling the end of the movement if u_{pex} is below threshold, represented through an inhibiting peak detector $u^{\mathrm{pd}}_{\mathrm{pex}}(t)$, and u_{pin} is above threshold, expressed with an excitatory peak detector $u^{\mathrm{pd}}_{\mathrm{pin}}(t)$. This results in the input

$$s^{\mathrm{CoS}}_{\mathrm{mov}}(t) = -c_{\mathrm{CoS,pex}}\sigma(u^{\mathrm{pd}}_{\mathrm{pex}}(t)) + c_{\mathrm{CoS,pin}}\sigma(u^{\mathrm{pd}}_{\mathrm{pin}}(t)). \tag{15}$$

The CoS node turns off the intention, which boosts u_{cur}, leading to an update of initial end-effector position in u_{ini}. Changing the initial position also alters the movement plan, which allows for corrective movements once the CoS is below threshold and the intention node reactivates. A CoD node monitors conditions through which the current internal movement fails. It listens to a mismatch detector field $u^{\mathrm{CoD}}_{\mathrm{mov}}(\boldsymbol{x}, t)$ comparing the current target position from the sensory stream with a working memory representation of target position at the start of the internal movement. In addition, it is also activated if $u_{\mathrm{ini}}(\boldsymbol{x}, t)$ does not contain a peak.

An outer loop of intention and CoS nodes $u^{\mathrm{int}}_{\mathrm{rea}}$ and $u^{\mathrm{CoS}}_{\mathrm{rea}}$ is activated through an external task input $s^{\mathrm{tsk}}_{\mathrm{rea}}(t)$ to reach to the target (see Figure 1, bottom). This can be compared to the go signal in Bullock, Cisek, and Grossberg [2]. In our architecture, however, this signal does not contribute to the shape of the velocity profile. Putting $s^{\mathrm{tsk}}_{\mathrm{rea}}(t)$ below threshold is the only way to suppress the otherwise fully autonomous movement generation. The CoS node compares initial end-effector position and the target position. Once these positions overlap, which is detected by a match detector field $u^{\mathrm{CoS}}_{\mathrm{rea}}(\boldsymbol{x}, t)$ receiving additive inputs from both u_{tar} and u_{ini}, the CoS of the reaching movement is fulfilled and the inner loop of $\boldsymbol{\lambda}$ movements and updates of initial positions is turned off. Note that the CoS node might deactivate once the target is moved, thus reactivating the inner loop and executing corrective movements.

F. The developmental stage of reaching

Three components of the architecture are probed for developmental effects by generating an "infant" version of each. Firstly, the lateral interactions in the DNF representing the initial end-effector position are assumed to be too weak to keep working memory of that position alive throughout the whole virtual movement, leading to a loss of the movement plan and thus a premature termination of the virtual movement. Secondly, the oscillatory layers may not include regions (and associated weights) large enough to generate large movement amplitudes. In other words, these fields are assumed to missing the weights necessary to reach the target in one cycle. The movement plan is thus mapped onto a region containing semi-random weights or to the outermost region that contains valid weights. Thirdly, the relation between planned Cartesian movement and muscle equilibrium points may not be fully known to the system (e.g., because stiffness and resting length control is not perfect). This is plausible for a growing body and developing muscle strength. This distorts the virtual end-effector movement, but also affects the corollary discharge pathway that estimates the internal end-effector position.

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

Impairing one or multiple of the components of the architecture listed above should have a significant influence on reaching behavior, leading to movements that feature multiple distinct movement units and a longer, less straight trajectory. Nevertheless, the autonomy of the architecture may bring the end-effector to the target location at some point. Sensory feedback about the achieved end-state may drive the learning process that reduces movement units and increase movement straightness over time. In this paper, we do not yet model this process of autonomous learning, however.

IV. Experiments

In this section, we will first evaluate a fully developed state of our architecture in which the mappings and weights have converged. Movement takes place in a 50 cm by 50 cm plane placed 20 cm in front of the robot and to the left of the robot's body center (see also Figure 1).

For all experiments, we use artificial visual inputs in form of fields of localized peaks of activation instead of real camera input to have full control of stimulus strength and position for reproducibility. We use the simulation solution Webots (http://www.cyberbotics.com) to execute the movements with the seven degrees-of-freedom arm. This ensures that the robot does not damage itself during execution of the movement commands using an impaired configuration of our architecture (generated movement might be jerky and unpredictable). The fully developed architecture was tested on hardware as well (RGB camera, Kuka arm), but this will not be discussed here.

A. Reaching movements and on-line updating

We first let the "adult" architecture reach for static targets in front of the robot. We vary starting position of the end-effector and target position, resulting in reaching movements in different directions and distances. The target positions are reached with a single virtual movement and subsequent movement of the end-effector. The velocity profiles of both virtual and external trajectories are bell-shaped (see Figure 2), with the virtual movement ending roughly at reaching peak velocity of the end-effector. Movement time is constant and does not depend on movement distance, which leads to a linear dependency between distance to target and peak velocity. Due to the transformation from Cartesian movement plan to joint space, the resulting trajectories are not perfectly straight.

We conduct the following experiment to test on-line updating in the "adult" architecture. We choose a two-step paradigm (see [33]) in which the end-effector starts in the center of an imaginary cross and the first target is placed on one of the four ends of the cross' equally long arms. During movement towards the first target, the target position switches, at varying inter-stimulus intervals (ISI), to the end of a neighboring cross arm. Sample trajectories for four different ISIs (600 ms, 700 ms, 800 ms, 900 ms) for this layout and one combination of targets are shown on the top left in Figure 3. Inspired by another experimental study of human on-line updating [8], we position the first target again on one of the arms of a cross, but then move the target perpendicular to the cross arm bearing

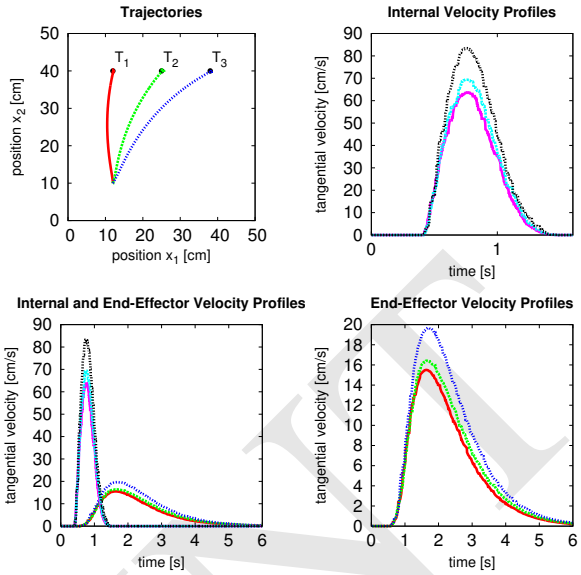

Fig. 2. Exemplary trajectories (top left) and profiles of tangential velocity for virtual movements (top right) and end-effector movements (bottom right) for different movement targets. The bottom left plot shows a combination of virtual and external profiles to show that the virtual movement ends roughly at peak velocity of the end-effector movement.

the target. The distance between first and second target is equal to the length of a cross arm. Sample trajectories of this second layout for the same four ISIs and for one combination of targets are shown on the top right in Figure 3. The resulting tangential velocity profiles feature two distinct movement units (see Figure 3, bottom row).

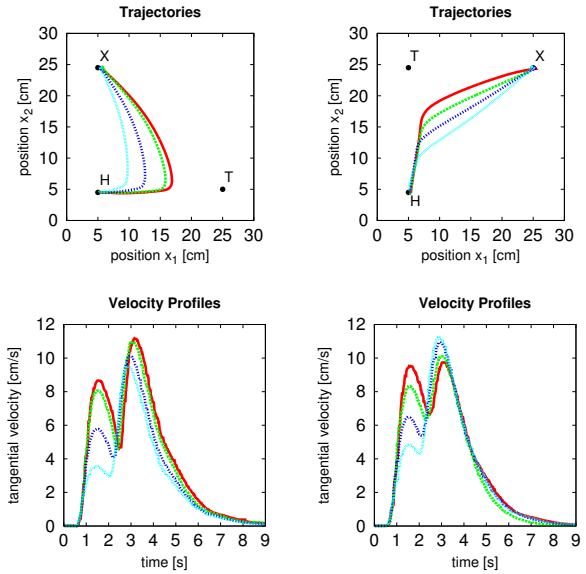

Fig. 3. Top row: Trajectories for different on-line updating setups (see text for details) and ISIs. The starting position of the hand is marked with the letter **H**, the first target position with **T** and the final target position with **X**. Bottom row: velocity profiles for the trajectories shown in the top row, displaying two movement units with varying peak velocities.

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

TABLE I
MEAN AND VARIANCE OF AMOUNT OF MOVEMENT UNITS AND STRAIGHTNESS 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

B. Reaching during the developmental stage

For the remaining experiments, we use the following setup. We define three starting positions S_1, S_2, S_3 for the robotic end-effector, which are placed on a line 30 cm in front of the robot along x_2 and 12, 25, and 38 cm to the left of the robot along x_1. Three targets T_1, T_2, T_3 are placed at the same position along x_1, but 30 cm further away in x_2.

We define three conditions of impairment: condition A is a lack of working memory sustain; condition B exhibits an underdeveloped movement timing system realized through incorrect oscillator weights for longer distances; condition C makes errors in translating the movement plan into joint space, due to a randomized offset added to the inverse kinematics. We let our architecture execute movements from the three starting positions to the three target positions under conditions A, B, C, and a combination of all of them (A+B+C). We record trajectories and velocity profiles. We then analyze the mean number of movement units (MUs) and straightness and compare these means with the fully developed architecture (which we call condition N). For each condition, we record two movements for each combination of start and target positions, summing up to 18 movements per condition.

Figure 4 shows sample trajectories and velocity profiles for movements from starting point S_1 to the three targets. For condition A, the resulting trajectories are straighter than for a fully developed architecture (see Figure 2). The velocity profiles are less smooth and feature multiple MUs. Condition B's trajectories are slightly less straight than condition N's. The velocity profiles show distinct MUs. Trajectories for condition C are less straight and exhibit overshoot. The velocity profiles contain few MUs with decreasing peak velocity. The combination of all three conditions shows the same characteristics of multiple MUs and a decrease in straightness. Mean amount of MUs and straightness over all combinations of starting positions and targets is listed in Table I. All impairment conditions show an increase in MUs. The straightness of Conditions A and B are comparable to the fully developed architecture, while Conditions C and A+B+C show a clear decrease in straightness. All impairment conditions exhibit an increase in movement time. Nevertheless, the target is successfully reached for all conditions despite the serious impairment of components.

V. CONCLUSION

We have presented a neuro-dynamic architecture for movement generation that captures both characteristics of adult goal-directed movement as well as infant movements around the onset of reaching. Autonomy is a key feature of the architecture as it ensures that movement to the target is achieved even when the target shifts during movement preparation or execution. We examined reaching movements generated by the architecture when three components were put into an earlier developmental stage (unstable working memory of initial position of the arm, gain of neural oscillator insufficient to reach target in single motion, inverse kinematics not correctly learned). We found two signatures of immature reaching, a larger number of movement units and curved end-effector paths.

A more complete picture of the development of reaching surely requires additional probes into the overall connectivity and autonomy of the movement generation architecture. We did not discuss here how the architecture produces different movement speeds at the same movement amplitude. The treatment of muscular control is also still rudimentary, lacking an account for co-contraction and the concomitant modulation of stiffness. The autonomous organization of the movement generation system may make it possible to produce more complex virtual trajectories, that may drive weak muscles more effectively.

ACKNOWLEDGMENT

The authors acknowledge the financial support of the European Union Seventh Framework Programme FP7-ICT-2009-6 under Grant Agreement no. 270247—NeuralDynamics. This work reflects only the authors' views; the EC is not liable for any use that may be made of the information contained herein. Discussion with Daniela Corbetta and Hendrik Reimann is gratefully acknowledged.

REFERENCES

[1] Christopher G Atkeson and John M Hollerbach. Kinematic features of unrestrained vertical arm movements. *The Journal of Neuroscience*, 5(9):2318–2330, 1985.

[2] Daniel Bullock, Paul Cisek, and Stephen Grossberg. Cortical networks for control of voluntary arm movements under variable force conditions. *Cerebral Cortex*, 8(1):48–62, 1998.

[3] Daniel Bullock and Stephen Grossberg. Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychological review*, 95(1):49–90, 1988.

[4] Mark M Churchland, John P Cunningham, Matthew T Kaufman, Justin D Foster, Paul Nuyujukian, Stephen I Ryu, and Krishna V Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, 2012.

[5] Daniela Corbetta, Sabrina L Thurman, Rebecca F Wiener, Yu Guan, and Joshua L Williams. Mapping the feel of the arm with the sight of the object: on the embodied origins of infant reaching. *Frontiers in psychology*, 5:576, 2014.

[6] Anatol G Feldman. Once more on the equilibrium-point hypothesis (λ model) for motor control. *Journal of motor behavior*, 18(1):17–54, 1986.

[7] J Randall Flanagan, David J Ostry, and Anatol G Feldman. Control of trajectory modifications in target-directed reaching. *Journal of motor behavior*, 25(3):140–152, 1993.

[8] Tamar Flash and Ealan Henis. Arm trajectory modifications during reaching towards visual targets. *Journal of cognitive Neuroscience*, 3(3):220–230, 1991.

[9] Apostolos P Georgopoulos, John F Kalaska, and Joe T Massey. Spatial trajectories and reaction times of aimed movements: effects of practice, uncertainty, and change in target location. *J Neurophysiol*, 46(4):725–743, 1981.

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015

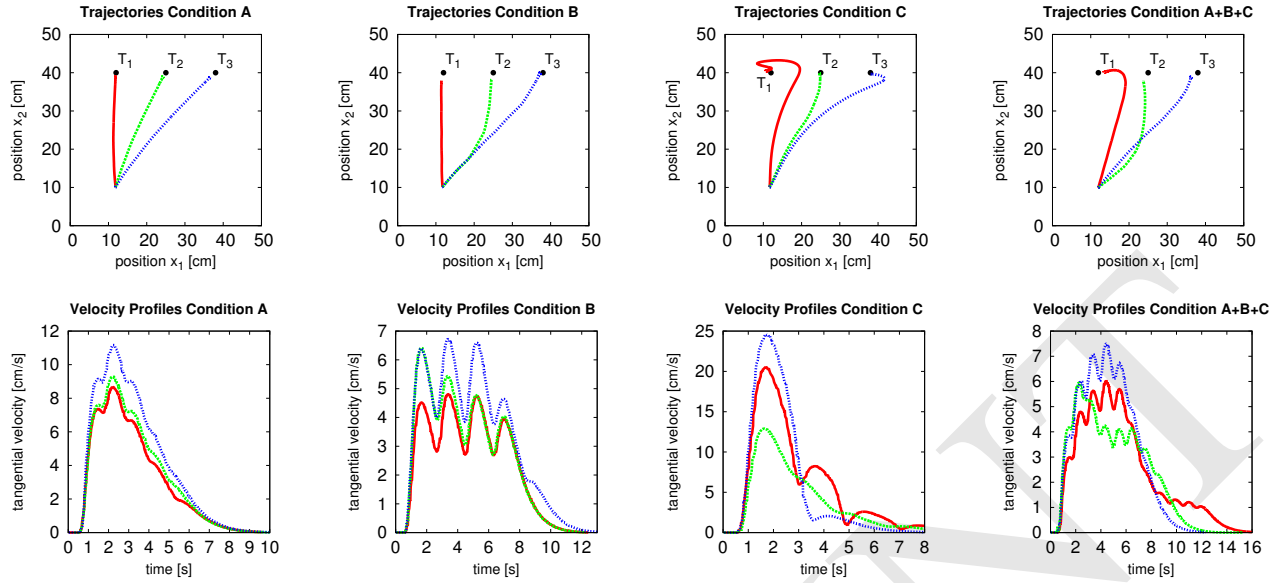

Fig. 4. This figure shows trajectories (top) and velocity profiles (bottom) for starting position S_1 to the three targets (solid red for T_1, bold-dotted green for T_2, dotted blue for T_3) and, from left to right, the conditions A, B, C, and the combination of all three.

[10] Paul L Gribble, David J Ostry, Vittorio Sanguineti, and Rafael Laboissière. Are complex control signals required for human arm movement? *Journal of Neurophysiology*, 79(3):1409–1424, 1998.

[11] Jeffrey S Johnson, John P Spencer, and Gregor Schöner. A layered neural architecture for the consolidation, maintenance, and updating of representations in visual working memory. *Brain research*, 1299:17–32, 2009.

[12] JA Scott Kelso. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. The MIT Press, 1995.

[13] Guido Knips, Stephan K U Zibner, Hendrik Reimann, Irina Popova, and Gregor Schöner. A neural dynamics architecture for grasping that integrates perception and movement generation and enables on-line updating. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014)*, pages 646–653, 2014.

[14] Franz Mechsner, Dirk Kerzel, Günther Knoblich, and Wolfgang Prinz. Perceptual basis of bimanual coordination. *Nature*, 414:69–73, 2001.

[15] Daniel W Moran and Andrew B Schwartz. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82(5):2676–2692, 1999.

[16] Pietro Morasso. Spatial control of arm movements. *Experimental brain research*, 42(2):223–227, 1981.

[17] Claude Prablanc and Olivier Martin. Automatic control during hand reaching at undetected two-dimensional target displacements. *Journal of Neurophysiology*, 67(2):455–469, 1992.

[18] Mathis Richter, Yulia Sandamirskaya, and Gregor Schöner. A robotic architecture for action selection and behavioral organization inspired by human cognition. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2457–2464, 2012.

[19] Uri Rokni and Haim Sompolinsky. How the brain generates movement. *Neural computation*, 24(2):289–331, 2012.

[20] Yulia Sandamirskaya and Tobias Storck. Learning to Look and Looking to Remember: A Neural-Dynamic Embodied Model for Generation of Saccadic Gaze Shifts and Memory Formation. In Petia Koprinkova-Hristova, Valeri Mladenov, and Nikola K Kasabov, editors, *Artificial Neural Networks SE - 9*, volume 4 of *Springer Series in Bio-/Neuroinformatics*, pages 175–200. Springer International Publishing, 2015.

[21] Gregor Schöner. A dynamic theory of coordination of discrete movement. *Biological Cybernetics*, 63:257–270, 1990.

[22] Gregor Schöner. Timing , Clocks , and Dynamical Systems. *Brain and Cognition*, 48:31–51, 2002.

[23] Gregor Schöner. Dynamical systems approaches to cognition. In Ron Sun, editor, *Cambridge Handbook of Computational Cognitive Modeling*, pages 101–126, Cambridge, UK, 2008. Cambridge University Press.

[24] Gregor Schöner and JA Scott Kelso. Dynamic pattern generation in behavioral and neural systems. *Science*, 239:1513–1520, 1988.

[25] Anne R Schutte and John P Spencer. Tests of the dynamic field theory and the spatial precision hypothesis: capturing a qualitative developmental transition in spatial working memory. *Journal of experimental psychology. Human perception and performance*, 35(6):1698–725, December 2009.

[26] Anne R Schutte, John P Spencer, and Gregor Schöner. Testing the Dynamic Field Theory: Working Memory for Locations Becomes More Spatially Precise Over Development. *Child Development*, 74:1393–1417, 2003.

[27] Vanessa R Simmering and Sammy Perone. Working memory capacity as a dynamic process. *Frontiers in psychology*, 3(January):567, January 2012.

[28] Soeren Strauss and Dietmar Heinke. A robotics-based approach to modeling of choice reaching experiments on visual attention. *Frontiers in psychology*, 3:105, 2012.

[29] Esther Thelen, Daniela Corbetta, and John P Spencer. Development of reaching during the first year: role of movement speed. *Journal of Experimental Psychology: Human Perception and Performance*, 22(5):1059–1076, 1996.

[30] Esther Thelen, Gregor Schöner, Christian Scheier, and Linda B Smith. The dynamics of embodiment: A field theory of infant perseverative reaching. *Brain and Behavioral Sciences*, 24:1–34, 2001.

[31] Emanuel Todorov and Michael I Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1236, 2002.

[32] Ya-Weng Tseng, John P Scholz, Gregor Schöner, and Lawrence Hotchkiss. Effect of accuracy constraint on joint coordination during pointing movements. *Experimental Brain Research*, 149(3):276–288, 2003.

[33] JF Van Sonderen, JJ Denier Van der Gon, and CCAM Gielen. Conditions determining early modification of motor programmes in response to changes in target location. *Experimental Brain Research*, 71(2):320–328, 1988.

[34] Claes von Hofsten. *Development of visually directed reaching: The approach phase*. Department of psychology, University of Uppsala [Psykologiska inst., Uppsala univ.], 1979.

[35] Stephan K U Zibner, Christian Faubel, Ioannis Iossifidis, and Gregor Schöner. Dynamic neural fields as building blocks of a cortex-inspired architecture for robotic scene representation. *Autonomous Mental Development, IEEE Transactions on*, 3(1):74–91, 2011.

Published as: Stephan K U Zibner, Jan Tekülve, and Gregor Schöner. The Neural Dynamics of Goal-Directed Arm Movements: A Developmental Perspective. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob), 2015 Joint IEEE International Conferences on*, pages 154–161, 2015