Increasing Autonomy of Learning Sensorimotor Transformations with Dynamic Neural Fields

Yulia Sandamirskaya* and Jörg Conradt†,
* Institut für Neuroinformatik, Ruhr-Universität Bochum
Universitätsstr. 150, 44780 Bochum, Germany
Email: yulia.sandamirskaya@ini.rub.de
† Institue of Autonomatic Control Engineering, TUM
Karlstr. 45, 80333 München, Germany
Email: conradt@tum.de

***Abstract*— In this paper, we introduce a neural-dynamic architecture that enables autonomous learning of sensory-motor mappings in a closed behavioral loop. Dynamic neural fields ensure stability of perceptual and motor representations, a neural-dynamic representation of the condition-of-satisfaction autonomously terminates the current action and enables activation of the next action, triggering a transient learning process when appropriate. Sampling of the high-dimensional space of the sensory-motor mapping is facilitated by a representation of the behaviorally salient states through localized activity peaks in the visual and motor neural-dynamic representations. We present the basic concepts of our autonomous learning architecture, a robotic implementation using a dynamic vision sensor mounted on a pan-tilt unit, which enables learning in a closed behavioral loop, and demonstrate functioning of the architecture in a simple learning scenario.**

I. INTRODUCTION

Sensorimotor transformations are at the core of behavior of autonomous cognitive agents, both robotic and human. The agent receives information about the state of the environment and its own body from sensors and has to generate motor commands based on this information in order to achieve its goals. The mapping from the perceived states to the motor commands – however complex and hierarchical it is – might change in dynamic real-world environments. Thus, adaptation of this mapping and its self-organization are necessary to enable autonomous functioning of the agent. Autonomy of the adaptation processes within sensory-motor mappings is challenged by the continuity in time and in space of the physical variables which define the environmental states. This autonomy relies on autonomy of the perceptual processes and motor control of the behaving agents, an issue often underdeveloped in work on learning sensory-motor mappings.

For instance, in the field of developmental robotics, motor bubbling is often used to discover the complex mapping between the sensory space and the motor space of the robot [1]. In an exploration phase, random movements are generated and the resulting sensory changes are recorded to drive learning process in a learning phase. The learning process itself applies gradient-descent optimization based on prediction and motor-control errors in a multi-layer perceptron. Here, a mapping between the sensory and motor spaces can be learned and the learning process may be optimized for efficiency of the exploration strategy and overall speed, but the problem of autonomy of the learning process is not addressed – there's no mechanism to autonomously detect when to initiate and to terminate an action, when to trigger and stop the learning process, errors are measured outside the learning system. Although such approach is useful to calibrate robotic systems, it does not provide insights into autonomous learning, as observed in humans and desired in truly autonomous agents [2]. The lack of autonomy in the coupling of the adaptive sensory-motor maps to the perceptual and motor systems is characteristics for developmental robotic architectures [3], [4], [5].

Other examples of algorithms for building and updating adaptive sensory-motor maps were introduced and applied in robotic scenarios [6], [7]. An architecture that combines a self-organizing map algorithm and dynamics of neural fields with reinforcement learning [7] exemplifies how a dynamical sensory-motor map may be formed, predictions of the upcoming sensory inputs may be generated, and sequences of states may be planned. The learning process here, however, consists of generating random movements, storing the respective commands and their consequences, and using the resulting data to drive the self-organization algorithm. The autonomy of this implementation is limited: important problems of autonomous segmentation and categorization of sensory inputs, as well as categorization of low-level motor commands, have to be solved in order to link the learning system to sensors and motors of a real physical agent. Similarly, in architectures in which the sensory-motor transformations are learned in a self-organizing map (SOM)[6], the learning process lacks autonomy. Here, the SOM builds a mapping between the sensory space of a simulated stereo camera and the end-effector of a simulated robotic arm, while actions are generated by sending random commands and observing sensory states when each action is finished. The mechanism of autonomous selection, initiation, monitoring, and termination of the actions are not included in the model. The moments in time, when it is appropriate to update the map are given to the algorithm, not detected autonomously. Most of the implementations of SOM algorithms do not focus on autonomy of the learning process [8].

Autonomy of cognitive processes and their development is central in the dynamical systems approach to modeling human cognition [9]. Dynamic field theory is a particular flavor of the dynamical systems approach, which has been particular successful in application of the cognitive models to control of robotic behavior [10], [11], [12]. The core element in this framework are dynamic neural fields (DNFs) – activation functions defined over topological spaces, which characterize the state of the behaving agent and the environment. Localized activity peaks emerge as stable solutions of the dynamics of DNFs and represent salient characteristics of the perceived states, as well as the goals of motor actions. Autonomy of representations in DFT is achieved by introduction of elements of intentionality [13], [14], which ensure activation and deactivation of the relevant representations when appropriate.

Here, we demonstrate how the framework of DFT can be applied to learning sensory-motor transformations and exploit the intentional structure of representations to enable autonomous learning, along with autonomous perception and action generation. The actions are initiated and terminated autonomously based on stabilized representations of the sensory inputs. The learning process is triggered autonomously when a match between the intended and the actual sensory state is detected and stabilized in the neural-dynamic representation of the condition-of-satisfaction.

We present here the first robotic implementation of the architecture in a looking scenario using a pan-tilt camera unit as the test-bed platform to explore autonomous learning of sensory-motor mappings. Looking behavior is unique in its simplicity and richness, with a vast room for learning and adaptation [15]. Here, we present the first stage of the autonomous adaptation of the neural-dynamic controller of this robot – adaptation of the mapping between the low-level sensory representation (a visual intention induced by a saliency map in retinal coordinates) and low-level motor-command representation (motor intention based on a map of desired pan-tilt configurations). The learning algorithm is similar to a simple reward-driven Hebbian learning rule, which strengthens projections between the sensory and the motor maps if a condition-of-satisfaction dynamics autonomously detects a match between the sensory input and the expected activation in the central part of this DNF when a successful saccade centers the selected object. Both the representation of the visually-induced intention and the respective motor command that result in a successful saccade, are stored by a working memory mechanism, inherent in the dynamics of neural fields. This memory mechanism sustains the relevant representations to enable learning.

The strength of this architecture is threefold: (1) we introduce stability to the sensory-motor representations, needed to drive the learning process, (2) we introduce the concept of condition-of-satisfaction in the sensory-motor maps learning, which enables autonomous activation and deactivation of the representations and drives the learning processes coupled to continuous sensory-motor dynamics, and (3) the architecture facilitates sampling of the space of the sensory-motor mapping through localized activity peaks in the neural fields dynamics. Next, we present the mathematical and conceptual basis of the model and details of the robotic implementation of the architecture.

II. THE MODEL

A. The basics of DFT

Dynamic Field Theory (DFT) is a conceptual and mathematical framework in which the neural dynamics accounts for emergence of elementary cognitive functions – such as detection of the most salient information, selection among alternatives, stabilization of decisions, and stabilization of working and long term memory – from the low-level, continuous dynamics of neural fields coupled to sensory input and to motor output [16]. The dynamic neural fields (DNFs) are activation functions defined over behaviorally relevant dimensions, such as visual features (e.g., color or intensity), spatial position of the target, or motor commands (e.g., the desired velocity or force). In DFT, the state parameters, such as the perceptual features or motor commands, are represented as stable localized activity bumps, or peaks. The positive activation within such peaks specifies values of the behavioral parameter that characterize the particular state. The peak-solution of a DNF is stabilized by lateral interactions within the neural field according to the neural field equation [17]:

$$\begin{aligned}\tau\dot{u}(x,t) = & - & u(x,t) + h + \int f(u(x',t))\omega(x'-x)dx' \\ & + & I(x,t) \end{aligned} \tag{1}$$

In Eq. 1, $u(x,t)$ is activation of the DNF, spanned over a behavioral dimension, x (e.g. visual space or motor commands), in time t; τ is a relaxation time-constant of the dynamics; h is a negative resting level ensuring that the DNF is salient if no input is present; $f(\cdot)$ is a sigmoidal non-linearity, which shapes the output of the neural field; ω is a Gaussian-shaped interaction kernel with the shot-range excitatory and the long-range inhibitory part; $I(x,t)$ is input to the DNF, coming either from a sensor or from another DNF.

The pattern of lateral interactions in a DNF ensures that localized peaks of supra threshold activation emerge as stable solutions of the Eq. 1. These peaks are units of representation in DFT and represent the percepts and

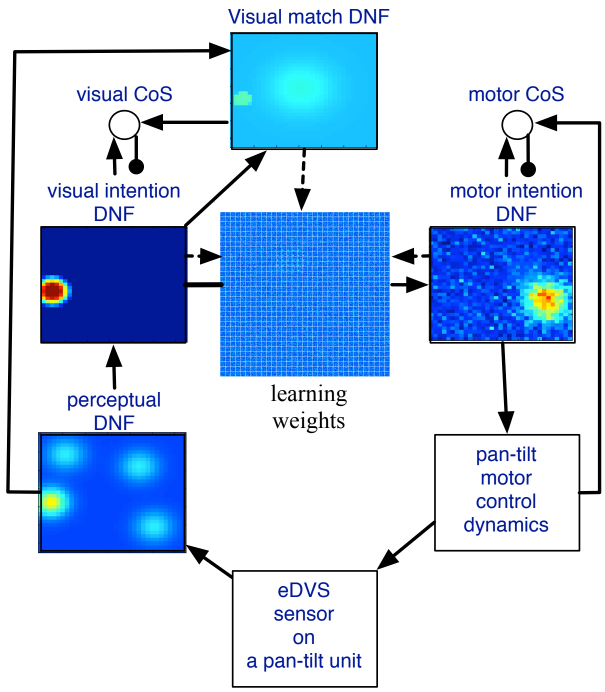

Fig. 1: The DNF sensory-motor map learning architecture. See text for details.

motor commands, which impact on the behavior of the agent.

B. Intentional structure

In order to enable autonomous activation and deactivation of states in a DNF architecture, each behaviorally relevant state has two dynamical components – an intention and a condition-of-satisfaction [11]. An intention is simply a dynamic neural field, Eq. 1, which is spanned over some behavioral space, receives input from the sensory surface, and is coupled to the down-stream structures of the architecture, ultimately setting attractors for low-level motor dynamics, which drive the behavior of the agent. Condition-of-satisfaction is another DNF, also following Eq. 1, which receives a subthreshold input from the intention DNF. This input is not sufficient to activate the CoS, but makes the CoS DNF more sensitive to certain inputs, i.e. the CoS DNF is preactivated, or preshaped, by the intention DNF. The preactivated CoS may be activated by a sensory input that matches the subthreshold input from the intention DNF. An active CoS field, in its turn, inhibits the intention DNF through a negatively weighted coupling. The motor action stops and the agent selects the next action according to a memorized sequence of required actions [14], rules of behavioral organization [11], or, in this work, driven by the most salient sensory input to the intention DNF.

In the architecture presented here, we use zero-dimensional CoS neural "fields", since the CoS of the motor system represents a match between the current motor state, defined by a scalar variable for each DoF – pan or tilt – of the robotic system, and the desired motor state, defined by the location of the activity peak in the motor DNF, also represented by a scalar variable for each DoF of the robot. The visual CoS is driven by a match DNF, which, additionally to its role as a CoS field also triggers the learning process, as described further (see Fig. 1).

C. The architecture

Fig. 1 shows the DNF architecture, which enables learning of the mapping between the visual representation of a salient object and the motor command needed to foveate this object, i.e. to bring the object in the central part of the visual field.

Here, the visual input from a dynamic vision sensor (DVS) [18] provides positive activation to the perceptual DNF, in which peaks of suprathreshold activation is built at locations, where salient pixels are concentrated. This perceptual system is a simple model of a saccade target selection system. The perceptual DNF is tuned in such a way that peaks are formed if a moving input is present in the visual array. When visual input ceases, peaks in the perceptual DNF decay and are not sustained; new moving input induces new peak(s) immediately.

The visual intention DNF, on the contrary, builds self-sustained activity peaks. A peak in this field represents the target for the upcoming saccade and has to be sustained for the time of the saccade. The visual intention peak may be switched off by two inputs: either by the visual CoS, which detects that the selected target is present in the central part of the image, or by the motor CoS, which signals that the saccade movement is accomplished. Both these inputs inhibit the visual intention field and bring it in the input-driven regime, in which the perceptual input may induce a new activity peak. Activation in the CoS structures ceases and inhibition on the intention DNF is released – the visual intention field may stabilize the representation of the new target.

Through the mapping between the visual and the motor DNFs, which is subject to the learning dynamics, a peak in the visual intention DNF induces a peak in the motor intention DNF. The motor intention DNF is coupled to the motor dynamics and drives looking behavior, which is modeled by a simple attractor dynamic that sets the pan and the tilt of the camera unit. Before the learning process has started, the mapping between the visual and the motor intention is not known and is modeled by a random map. A peak at an arbitrary location in the motor intention DNF is built from the random input, induced by the lateral interactions in this field. This peak generates a looking action, which, initially, does not lead to centering the object. In this case, the motor CoS is activated (the movement of the amplitude encoded by the motor intention is performed) and the visual and motor intention fields are inhibited to be input-driven again; the self-sustained activity peaks decay and give way to the new perceptually-driven activity peaks. The

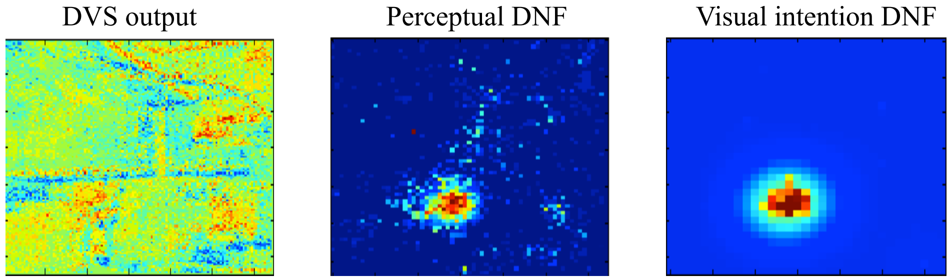

Fig. 2: Selection of the saccade target on the sensory pathway from the DVS camera (left), through perceptual DNF (middle), to visual intention DNF (right).

visual CoS is not activated in this case and no learning happens.

When after a looking act, the target appears in the central part of the visual field, the visual match DNF is activated and triggers a learning process, in which the (still active) location in the visual intention DNF and the (still active) location in the motor intention DNF are associated by strengthening the respective locations in the sensory-motor map, according to a simple Hebbian-like ("fire together – wire together") learning rule, Eq. 2:

$$\tau_l \dot{T}(x,y,k,l) = \lambda \int f(u_{match}(x,y))dxdy \cdot$$
$$\cdot \Big(-T(x,y,k,l) + f(u_{vis}(x,y)) \times f(u_{mot}(k,l)) \Big) \quad (2)$$

Here, the mapping $T(x,y,k,l)$ (time-dependence is omitted in the equation) between the visual intention field, $u_{vis}(x,y)$, defined over image coordinates (x,y), and the motor intention field, $u_{mot}(k,l)$, defined over the motor coordinates, k (pan) and l (tilt), retains its values if the match DNF, $u_{match}(x,y)$, is salient. If there's a positive activation in the match DNF, the integral before the learning term shunts the change in the mapping to be non-zero. The learning equation sets an attractor for $T(x,y,k,l)$ at the values of positive correlation between the two intention DNFs, calculated as a sum between the output of the visual intention field, expanded along the dimensions of the motor intention field, and the output of the motor intention field, expanded in the dimensions of the visual intention field, augmented with a sigmoidal threshold function. This correlation computation is equivalent to calculations performed by gain-field cells and is one of the cognitive operations, which may be employed in the DFT [10].

D. Sensory-motor system

We have coupled the DNF architecture to a custom-made sensory-motor system, which consists of a dynamic vision sensor (DVS) – a retina-inspired neuromorphic camera, which detects change events in its sensor array and sends the information about the pixel and time of occurrence of each event asynchronously to the computer, on which the DNF architecture is simulated. A snapshot of a typical output of the DVS is shown in Fig. 2. The perceptual DNF is activated in a region of the most salient and spatially coherent activity in the DVS output. Because of the weak lateral interaction in the perceptual DNF, several activity peaks may be built in this field. The perceptual DNF is coupled to the visual intention DNF, which, in its turn, has strong lateral interactions, and consequently a single activity peak is built in this field over one of the salient locations. This peak represents the visual target of the next planned saccade.

Activity in the visual DNF is propagated to the motor DNF through the mapping, which is random in the beginning of the experiment and results in a peak in the motor intention DNF at a random location – the motor intention DNF performs a selection decision and stabilizes this decision, induced by the random input from the sensory-motor mapping. A peak in the motor intention DNF sets an attractor for the pan and the tilt of the robotic system (Fig. 3), the motor dynamics drives the camera towards the corresponding pose. Activity peaks in both visual and motor intention DNFs stay in the same locations during camera motion because of the strong lateral interactions in these fields, sufficient to induce self-sustained activity, independent of the changing inputs. This activity ceases when motor CoS is activated, which brings both fields back in the input-driven regime.

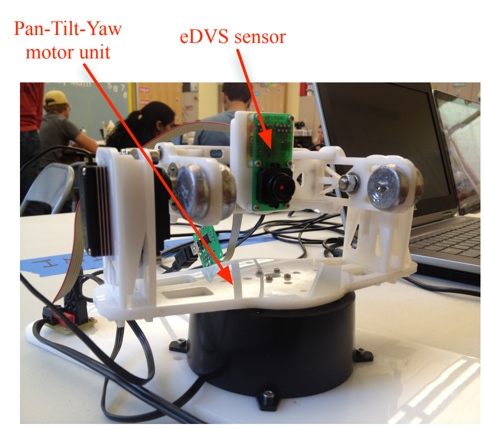

Fig. 3: The DVS sensor mounted on a pan-tilt motor unit.

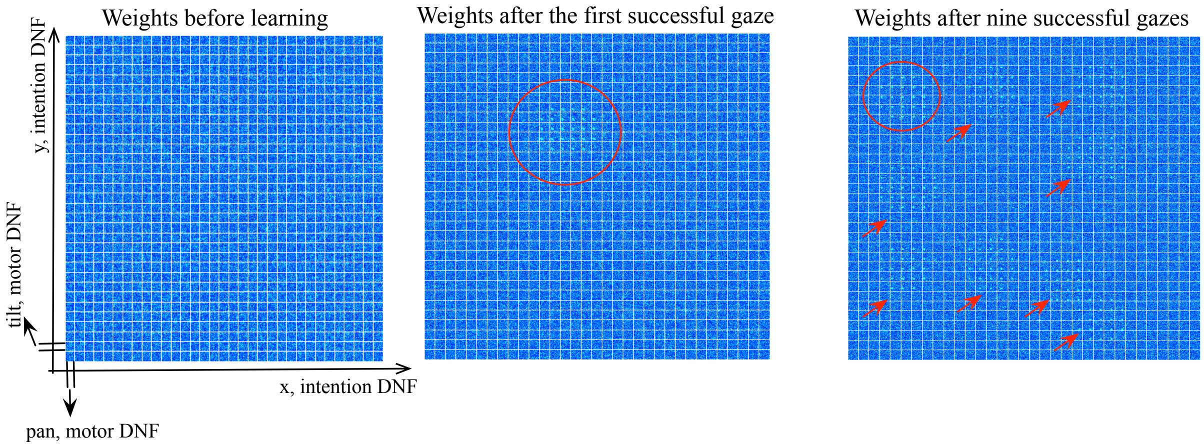

Fig. 4: The mapping between the visual and the motor intention spaces. The colormap of the plots spans the interval [0, 1] (light blue to dark blue). The 4D mapping is shown here as slices along the motor dimensions, arranged in the figure according to the visual dimensions. Before learning (left), the mapping is initialized as random connections tensor. After the first successful saccade (middle), one region in the 4D field encodes a learned mapping between a localized region in the visual intention space and a localized region in the motor intention space (the region marked with the red circle; note the light-blue dots in the tiles in this region). After only a few successful saccades (nine shown here), a large portion of the 4D space of the mapping is learned (regions marked by the red circle and the red arrows).

E. Scenario and experimental set-up

In our experiment, the robot is put in front of a screen, on which a pattern of blinking inputs is present, which induces an activity pattern in the DVS output. The perceptual DNF stabilizes this output and the visual intention DNF selects one of the active regions as the target of the following saccade. Through the randomly initiated mapping, the output of the visual intention DNF provides a randomly distributed input to the motor intention DNF. Because of the lateral interaction in the motor intention DNF, a localized activity peak is built in this DNF at one of the locations. The robot performs a saccadic movement (controlled by a simple attractor dynamics), which stops when the motor attractor is reached. The motor CoS is activated and inhibits the intention DNFs slightly. If the saccade is not successful, the visual intention DNF builds peak over the new location of the selected visual object. The new visually-induced intention is again propagated to the motor intention DNF, where a new activity peak represents a new hypothesis about the required motor action.

This process is repeated until, by chance, one of the saccades brings the visual input close enough to the central portion of the visual array. This event induces an activity peak in the match DNF, which has a sub-threshold preactivation in its central part, projected from the active visual intention DNF, and is sensitive to visual input in the central part of the camera image. The active match DNF triggers the learning process in the mapping between the visual and the motor intention DNFs (according to the Eq. 2), while inhibiting the intention DNF additionally. Learning happens in a transient, when both intentions that have resulted in the successful saccade are still active. When the motor intention DNF is inhibited below threshold, peak in the match DNF decays as well, and the update of the sensory-motor weights stops. The new visual input from the periphery of the camera image induces a new activity peak in the visual intention DNF field, and the exploration and learning processes continue autonomously.

III. RESULTS

Fig. 4 shows the sensory-motor map before the learning session, after the first successful saccade, and after several successful saccades. Although the mapping is defined between two continuous spaces, both these spaces are sampled by bell-shaped localized activity peaks in the learning process. This facilitates the learning process, since the mapping is learned in a localized, finite, region of the 4D mapping space each time a successful saccade is made. Thus, the whole mapping space maybe sampled in as few as twenty successful saccades. By adjusting the width of the central region, preactivated in the match DNF, learning may be further refined. Here, we present only a proof of concept implementation of the autonomous learning architecture and demonstrate how learning of a sensory-motor mapping may be performed autonomously, in a closed behavioral loop. A thorough statistical analysis of the learning process is currently being performed.

IV. DISCUSSION

In this paper, we have presented a neural-dynamics architecture that enables autonomous learning of a sensory-motor mapping between the visual and motor intentions, each represented by a dynamic neural field. We have implemented this architecture on a robotic agent and demonstrated how learning accompanies autonomous generation of looking actions based on low-level sensory input in a closed behavioral loop. We have combined stability of the dynamic neural field representations with elements of the behavioral organization – the intention and condition-of-satisfaction neural-dynamical structures – to enable autonomy of the learning process, which includes autonomy of selection of the visual target, initiation of the motor action, termination of the motor action, and decision to trigger the learning dynamics. All these processes are controlled by autonomously generated instabilities in the dynamics of the neural fields, which are mutually coupled and linked to the raw, continuous in time and in space, sensory input. We have demonstrated autonomy of this learning architecture in a challenging setting with a dynamic vision sensor and an uncalibrated pan-tilt unit. Starting with a randomly generated coupling, the robot was able after several successful saccades to coarsely sample the sensory-motor mapping between the image-based visual intention and the map of the pan-tilt motor poses, which bring the target object into the center of the robotic camera. This first implementation demonstrates how learning may be accomplished in a closed behavioral loop and will be further analyzed and developed in subsequent experiments.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the financial support of DFG SPP *Autonomous Learning*, within Priority Program 1567; project "Development as autonomous learning: Emergence of developmental stages that lead from sensori-motor behaviors to embodied cognition".

References

[1] R. Saegusa, G. Metta, G. Sandini, and S. Sakka, "Active motor babbling for sensorimotor learning," in *Robotics and Biomimetics, 2008. ROBIO 2008. IEEE International Conference on*, pp. 794–799, IEEE, 2009.

[2] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan, "Principles of sensorimotor learning.," *Nat Rev Neurosci*, vol. 12, pp. 739–51, 12 2011.

[3] C. Gaskett and G. Cheng, "Online learning of a motor map for humanoid robot reaching," 2003.

[4] M. Kuperstein, "Infant neural controller for adaptive sensory-motor coordination," *Neural Networks*, vol. 4, no. 2, pp. 131–145, 1991.

[5] G. Metta, G. Sandini, L. Natale, and F. Panerai, "Development and robotics," in *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*, pp. 33–42, 2001.

[6] H. J. Ritter, T. M. Martinetz, and K. J. Schulten, "Topology-conserving maps for learning visuo-motor-coordination," *Neural networks*, vol. 2, no. 3, pp. 159–168, 1989.

[7] M. Toussaint, "A sensorimotor map: Modulating lateral interactions for anticipation and planning," *Neural Comput.*, vol. 18, pp. 1132–1155, 5 2006.

[8] R. Schulz and J. A. Reggia, "Temporally asymmetric learning supports sequence processing in multi-winner self-organizing maps," *Neural Computation*, vol. 16, no. 3, pp. 535–561, 2004.

[9] E. Thelen and L. B. Smith, *A Dynamic Systems Approach to the Development of Cognition and Action.* Cambridge, Massachusetts: The MIT Press, A Bradford Book, 1994.

[10] Y. Sandamirskaya, S. Zibner, S. Schneegans, and G. Schöner, "Using dynamic field theory to extend the embodiment stance toward higher cognition," *New Ideas in Psychology. Special Issue "Adaptive Behavior".*, in press.

[11] M. Richter, Y. Sandamirskaya, and G. 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*, 2012.

[12] S. K. U. Zibner, C. Faubel, I. Iossifidis, and G. Schöner, "Dynamic neural fields as building blocks for a cortex-inspired architecture of robotic scene representation," *IEEE Transactions on Autonomous Mental Development*, vol. 3, no. 1, pp. 74–91, 2011.

[13] J. R. Searle, *Intentionality — An essay in the philosophy of mind.* Cambridge University Press, 1983.

[14] Y. Sandamirskaya and G. Schöner, "An embodied account of serial order: How instabilities drive sequence generation," *Neural Networks*, vol. 23, pp. 1164–1179, December 2010.

[15] J. J. Hopp and A. F. Fuchs, "The characteristics and neuronal substrate of saccadic eye movement plasticity.," *Prog Neurobiol*, vol. 72, pp. 27–53, 1 2004.

[16] G. Schöner, *Cambridge Handbook of Computational Cognitive Modeling*, ch. Dynamical systems approaches to cognition, pp. 101–126. R. Sun, UK: Cambridge University Press, 2008.

[17] S. Amari, "Dynamics of pattern formation in lateral-inhibition type neural fields," *Biological Cybernetics*, vol. 27, pp. 77–87, 1977.

[18] J. Conradt, R. Berner, M. Cook, and T. Delbruck, "An embedded aer dynamic vision sensor for low-latency pole balancing," in *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pp. 780–785, IEEE, 2009.