


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.


In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

An embodied account of serial order: How instabilities drive sequence generation

Yulia Sandamirskaya*, Gregor Schöner¹

Institut für Neuroinformatik, Ruhr-Universität Bochum, Universitätsstr. 150, D-44780 Bochum, Germany

ARTICLE INFO

Article history:

Received 28 May 2010

Received in revised form 29 July 2010

Accepted 30 July 2010

Keywords:

Dynamic Field Theory

Sequence generation

Embodiment

Neural dynamics

Attractors and instabilities

ABSTRACT

Learning and generating serially ordered sequences of actions is a core component of cognition both in organisms and in artificial cognitive systems. When these systems are embodied and situated in partially unknown environments, specific constraints arise for any neural mechanism of sequence generation. In particular, sequential action must resist fluctuating sensory information and be capable of generating sequences in which the individual actions may vary unpredictably in duration. We provide a solution to this problem within the framework of Dynamic Field Theory by proposing an architecture in which dynamic neural networks create stable states at each stage of a sequence. These neural attractors are destabilized in a cascade of bifurcations triggered by a neural representation of a condition of satisfaction for each action. We implement the architecture on a robotic vehicle in a color search task, demonstrating both sequence learning and sequence generation on the basis of low-level sensory information.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

Even the simplest activities in life involve generating well-ordered sequences of actions. Preparing a meal and setting the table, packing the briefcase and driving to the office, or even just getting up from your desk to refill your glass at the water cooler are examples. Playing music, writing, and gesturing all entail serially ordered sequences of actions. Sequences are also the basis for the highest forms of cognition such as generating and comprehending language. Generating such action sequences involves many different processes including creating, working and long-term memories for the elements of the sequence, initiating any individual action while inhibiting other available actions, controlling and timing elementary actions, terminating a completed action and selecting the next action.

Naturally, how sequences are generated has been a central topic of psychology since Lashley's seminal insight that the serial order in a sequence is a separate and critical dimension of such behavior (Lashley, 1951). This insight was supported by the observation of characteristic patterns of errors in serial order tasks as well as by how response times depend on serial position, sequence length, and other factors. Modern theoretical work on sequence generation has addressed aspects of serial order across a range of tasks including speech production (Dell, Chang, & Griffin, 1997; Hartley & Houghton, 1996), spelling (Glasspool & Houghton,

2005), action planning (Cooper & Shallice, 2000; Grossberg, 1978), immediate serial recall (Henson, 1998), and free recall (Farrell & Lewandowsky, 2002). In many cases, sequences have been studied and modeled in relatively disembodied form, in which item and order memory is probed through simple, invariant motor responses such as keyboarding or button pressing. Such sequences are typically highly constrained. For instance, playing a piece of piano music entails sequential key presses, the order and duration of which is predetermined by the score. What little variance remains is constrained by speed–accuracy trade-offs (Pfordresher, Palmer, & Jungers, 2007).

In this paper we look at embodied systems that are situated in real environments and must make variable movements oriented at physical objects to achieve the goal of each sequence element. Setting a table, for instance, is a task that requires that sequences of individual actions be performed, each directed at an object and each potentially taking different amounts of time. This example highlights five critical properties of the embodied sequence generation.

First, embodied sequence generation is *autonomous*, that is, decisions to bring a particular object into the foreground, select it as the target of an action, and initiate that action are all driven by intrinsic processes that are coupled to the embodied system's own sensory and motor surfaces. As you set the table, nobody triggers your actions from the outside.

Second, each action within an embodied sequence is *part of a motor and perceptual continuum*. The particular movements required to set the table will vary each time depending on how you approach the table and where the dishes are stored. The precise grasp required to pick up a cup will vary with the pose of the cup which must be estimated from its visual appearance. Any

* Corresponding author. Tel.: +49 234 3227996; fax: +49 234 3214209.

E-mail addresses: yulia.sandamirskaya@ini.rub.de, sandayci@rub.de (Y. Sandamirskaya), gregor.schoener@ini.rub.de (G. Schöner).¹ Tel.: +49 234 3227965; fax: +49 234 3214209.

stage of the behavioral sequence thus really entails a continuum of possible percepts and movements out of which a specific instance is generated in the environmental context on each occasion. Although the perceptual objects and actions relevant to each stage of the sequence may be described as belonging to a particular category, their categorical representation alone is too poor to explain the observed behavioral flexibility when context varies.

Third, embodied sequence generation is *flexibly timed*. Even though each element in the table-setting sequence may take varying amounts of time depending, for example, on how far you need to reach to set down a piece, each action is brought to a conclusion before the next one is initiated. This is true also if an action suffers perturbations, for instance, if you do not find the particular dish you were intending to pick up next. Embodied sequence generation must be stable against variations of the time needed to terminate each action and must be able to use sensory information to determine that a subgoal of the sequence has been reached.

Two additional constraints arise from the demand that an understanding of embodied sequence generation be based on neuronal principles. Neural processing occurs *continuously in time*, sampled by asynchronous neuronal activity and continuously coupled to sensory inputs. Neuronal processing does not consist of discrete computational steps from one action or perceptual state to the next. Because neuronal time is fundamentally continuous, the emergence of discrete transitions between different sequence elements is in need of explanation.

Finally, neuronal representations naturally capture the continua of possible actions and possible sensory states through *graded representations*. That fact that neurons are discrete entities does not lead to observable discreteness of behavior. This is the basis for the neuronal concept of population coding (Deadwyler & Hampson, 1995; Erickson, 1974; Georgopoulos, 1991). Conversely, there are behavioral signatures suggesting that the metric dimensions of perceptual and motor representations are encoded. For instance, drift of metric memory over delays and the dependence of performance on task metrics can be accounted for using population coding ideas (Johnson, Spencer, & Schöner, 2008). How the categorical states that characterize the different stages of a reproducible sequence emerge from such underlying continua is thus in need of explanation.

Because current neuronal models of sequence generation (Botvinick & Plaut, 2006; Brown, Preece, & Hulme, 2000; Burgess & Hitch, 1999; Deco & Rolls, 2005; Elman, 1990; Houghton, 1990; Page & Norris, 1998) address these constraints only partially, there is, to our knowledge, no comprehensive theoretical account for embodied and situated sequence generation that explains how behavioral sequences are acquired and produced in the real world.

Our goal in this paper is to establish such an account. By employing the framework of Dynamic Field Theory (DFT) we ensure that this approach is neuronally based and consistent with the two constraints of *time-continuous* processing on *graded* representations (Schöner, 2008). The units of representation in dynamic fields are peaks of activation defined over continuous behavioral dimensions. Such peaks emerge *autonomously* from the interplay of neuronal interaction and sensory input and may be localized as dictated by current input or also be determined by memory traces acquired during learning. As a result, dynamic fields are capable of both representing categorical as well as *continuous metric* information. Because the peaks are attractor states of the underlying neuronal dynamics, they may be sustained over unpredictable delays enabling *flexibly timed* actions.

The critical element of our account is how transitions from one action to its successor in a sequence are generated. Sensory information signals that the condition of satisfaction for the current action is met. An instability emerges from the interaction of a neural

field representing actions with the neural field representing their conditions of satisfaction. That instability drives a transition between sequential actions.

The serial order of individual actions is represented in our account through a network of dynamic neural nodes. The discrete units in this network project onto continuous dynamic fields that represent the potential continua of actions and perceptual states. These perceptual and action representations are linked to the sensors of the agent, enabling it to learn and monitor sequences, and to the effector systems of the agent enabling it to generate context-appropriate actions. The network of dynamic neural nodes representing serial order is in this sense separate from the sensorimotor representations. This conception is in agreement with psychological (Henson, 1998) and neurophysiological (Aldridge & Berridge, 1998; Procyk, Tanaka, & Joseph, 2000) signatures of sequential action.

We demonstrate how this architecture can generate sequential behavior of an embodied agent by implementing the dynamics on a robot that performs a simple sequential task. In our implementation, sensory input comes from a camera and photoelectronic sensors. Sensory information is relatively low-level (e.g., unsegmented color-space distributions) illustrating how action sequences can be acquired and produced based on graded, noisy sensory information from real sensors. The sequence is performed by generating simple movements (driving around and approaching colored targets) on a robot that operates in an *a priori* unknown environment purely based on its own sensory information. The real-world implementation of sequence generation under these constraints provides a sufficiency proof for the embodied nature of our account.

After briefly reviewing the conceptual and mathematical framework of DFT in Section 2, we give a conceptual overview of the DFT sequencing architecture in Section 3, and a more thorough mathematical description in Section 4. An exemplary implementation of the architecture on a miniature robotic vehicle is described in Section 5, followed in Section 6 by the results from a series of robotic demonstrations. The discussion in Section 7 positions our approach within the theoretical literature and reviews its relationship to the relevant neurophysiology.

2. DFT as a framework for embodied sequence generation

Dynamic Field Theory (DFT) is a phenomenological description of the dynamics of neuronal activation that abstracts from the discreteness of neurons as units of neuronal networks and from a biophysical mechanism of spiking (Amari, 1977; Ermentrout, 1998; Wilson & Cowan, 1973). As a theoretical language, DFT has been widely used to model motor, perceptual, and cognitive functions (Schöner, 2008). The central idea of DFT is that neural representations are defined over continuous dimensions such as the direction and amplitude of a reaching movement or the retinal location or color of a visual stimulus. These dimensions span a conceptual motor or feature space, over which an activation function is defined. For each location, x , along these dimensions, the activation function, $u(x)$, indicates the presence of information by high levels of activation and the absence of information by low levels of activation. High levels of activation imply that neuronal representations or effector systems onto which the activation field projects are impacted by the activation variable. Low levels imply that no such impact takes place. The activation function, $u(x, t)$, evolves in time, t , as described by the integro-differential equation (Amari, 1977):

$$\tau \dot{u}(x, t) = -u(x, t) + h + \int f(u(x', t)) \omega(x - x') dx' + S(x, t). \quad (1)$$

Here, τ is a time constant of the dynamics that determines how

quickly the activation function, $u(x, t)$, relaxes to an attractor state that emerges from the stabilization factor, $-u(x, t)$, and the additive contributions: the negative resting level, $h < 0$, the lateral interactions shaped by the kernel, $\omega(x - x')$, and the external input, $S(x, t)$. The kernel is a bell-shaped function containing both excitatory connectivity of strength c_{exc} over the range, σ , and inhibitory connectivity of strength c_{inh} over all ranges (2). The convolution is implemented using periodic boundary conditions.

$$\omega(x - x') = c_{exc} \exp\left[-\frac{(x - x')^2}{2\sigma^2}\right] - c_{inh}. \quad (2)$$

The homogeneity of the connectivity pattern is characteristic for dynamical neural fields and, together with the bell-shaped kernel stabilizes localized peaks of activation as attractor solutions. This requires a non-linearity, $f(u(x, t))$, that thresholds the output of the field's dynamics as a soft sigmoid:

$$f(u(x, t)) = \frac{1}{1 + \exp[-\beta x]}. \quad (3)$$

The field dynamics has different dynamic regimes. In some instances, these regimes can be determined based on analytical techniques (Amari, 1977; Ermentrout, 1998; Wennekers, 2002), which is an advantage of this particular mathematical formulation of a neural dynamics and one of our motivations for adopting this form. In the absence of external input, one attractor state has activity in the field close to the negative resting level, h . This sub-threshold solution remains stable when weak localized input, $S(x, t)$, is introduced as long as the summed activation level, $h + S(x, t)$, does not surpass levels at which the lateral interaction becomes engaged. When that threshold is passed, the output, $f(u(x, t))$, and the interaction function drive the system through the *detection* instability, in which the sub-threshold solution disappears. Activation grows near the field sites at which localized external input was largest, developing into a localized peak of activation that inhibits the activation field elsewhere. The peak solution is stabilized against decay by the local excitatory interaction and against lateral diffusion by global inhibitory interaction. Such peaks thus form the second category of attractor states, which are self-stabilized by intra-field interactions, but also track changing localized input. Self-stabilized, localized activation peaks are the units of representation in DFT. For instance, objects in the visual field may be represented as peaks in a field spanned by visual features and visual space. Action plans or motor primitives may be represented by peaks in fields spanned by movement parameters. In DFT, all motor, perceptual, or cognitive states are represented by such attractor states, so that they are resistant to variations in input or coupling across different fields.

The sub-threshold solution may be inhomogeneous with various localized small bumps of activation due to weak input that represents learned categories within a feature dimension. Such inhomogeneities may be amplified into self-stabilized peaks by pushing the field through the detection instability by an increase of the resting level that provides a “boost” to the overall level of activation.

When localized input is removed, the detection instability may be experienced in reverse. In that *forgetting* instability, the self-stabilized peak becomes unstable and the system relaxes to the sub-threshold attractor. The forgetting instability happens at lower levels of localized input than the detection instability, so that there is a bistable regime in which detection decisions are stabilized. For sufficiently large resting levels, h , or strong lateral excitation in the neural field, the forgetting instability may not occur even when localized input is removed entirely. In this case, the self-stabilized peak is sustained in the absence of the localized input that first induced it. The forgetting instability may be induced, however, by lowering the resting level.

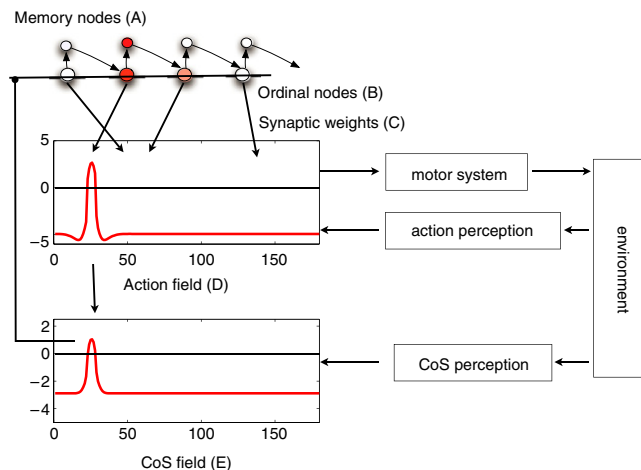


Fig. 1. Functional modules of the DFT sequence generation architecture.

The two instabilities translate graded changes in a neural field's input into macroscopic decisions represented by bifurcations in the field dynamics. These decisions are elementary forms of embodied cognitive function. Through these instabilities, dynamic fields provide an interface between the sensory and motor surfaces and the associated sensory-motor processes on the one hand, and a cognitive dynamics on the other hand, that enables agents to pursue goals that are not at all times dependent on sensory input.

3. The DFT sequence generation architecture

Modeling serially ordered sequence generation within the framework of DFT poses two problems. First, the system must represent at any time where it is within a sequence that is either being learned or being generated. Fields may represent spatial patterns of activation. Representing such an ordinal position requires, however, that activation patterns be generated sequentially in time. Second, achieving the temporal transitions from one activation pattern to the next requires solving a conflict with the stability of each of the activation patterns. The prior state in a sequence must first be destabilized in order to bring about the next state. Moreover, that instability must be induced reliably by sensory signals that indicate that the ongoing action has been accomplished. Each stage of a sequence must, therefore, be represented jointly with its termination condition.

To solve these problems, we introduce an architecture which combines neural field representations of different actions with a neuronal dynamics of discrete nodes that represent ordinal positions in a sequence of actions. We also introduce dynamic machinery to trigger a chain of instabilities in the dynamics of neural fields and discrete nodes that enables the transition from one action to the next in the sequence. The instabilities can be induced by a sensory signal or by an internally generated event. An overview of the DFT sequence generation architecture is given in Fig. 1.

First, we step through the main functional components of the architecture. A set of discrete neural-dynamic nodes represents ordinal positions within a sequence at which the system is at any point in time (Fig. 1(B)). This set of *ordinal nodes* is inspired by neurophysiological evidence for a neuronal representation of serial position within a sequence of actions. Neurons responsive to ordinal information have been found in neostriatal cortex (Aldridge & Berridge, 1998), in motor cortex (Carpenter, Georgopoulos, & Pellizzer, 1999), and in the supplementary motor area (Clower & Garrett, 1998). The ordinal nodes in our architecture have bistable activation dynamics that share inhibition so that only

one ordinal node can be active at a time. Asymmetrical excitatory connections between the nodes control the direction in which activation spreads from the beginning to the end of a sequence. These directed connections are mediated by *memory nodes* (Fig. 1(A)) that stay active during a transition phase between two successive actions.

When an ordinal node is activated, its projection onto an action field becomes effective. One or multiple *action fields* (Fig. 1(D)) span the space of actions that are possible at any given stage of the sequence. Which action is associated with a given stage is specified by the *synaptic connections* (Fig. 1(C)) from the corresponding ordinal node to the action field. During sequence learning, those connections are strengthened that project onto the field sites that are activated by the demonstrated action. The set of synaptic weights thus constitutes the memory for a graded representation of the action associated with a particular ordinal position. As an ordinal node becomes active, it induces a self-stabilized peak of activation at the location, to which the ordinal node projects. Such a peak in an action field then elicits motor behavior. The *motor system* is a dynamics that is controlled by the movement parameter values, over which the peak in the action field is centered. This dynamics couples to the effector system of the agent. The action field may receive perceptual input from the environment, while a particular action is being executed or observed. Through this perceptual channel, peaks of activation may be induced in the action field during sequence learning in response to a demonstration of a target action. During sequence production, a lower resting level makes that perceptual input alone not sufficient to induce an activation peak. The *action–perception* dynamics controls how this perceptual input is extracted from the sensory surface.

To solve the problem of destabilizing the stable state that specifies the current action, we explicitly represent the fact that this action has achieved its goal. This is inspired by the concept of a condition of satisfaction in Searle's analysis of intentional acts (Searle, 1983). The third component of the architecture is, therefore, a *condition of satisfaction (CoS) field* defined over a relevant perceptual dimension specific to the space of possible actions (Fig. 1(E)). The CoS field receives localized input from the active action field. This input pre-activates those locations in the CoS field that encode perceptual states indicating that the specified action has terminated. Such pre-activation makes the CoS field more responsive to input consistent with those perceptual states. Sensory input signaling the potential termination of actions is obtained from the *CoS-perception* dynamics. Where that sensory input overlaps along the field dimension with the localized pre-activation induced by the action field, the detection threshold is reached and a self-stabilized peak is formed in a detection instability. That peak inhibits the ordinal set and triggers a cascade of instabilities in the different sub-systems of the architecture with the end-result that the next ordinal node is activated and the system proceeds to executing the next action.

4. Mathematical structure of the DFT architecture of sequence generation

4.1. Ordinal nodes

The ordinal system consists of ordinal nodes with activities, d_i , and associated memory nodes with activities, d_i^m , where the index, i , counts the actions. The activity of these nodes evolves in time according to a set of coupled differential equations:

$$\begin{aligned} \tau \dot{d}_i(t) = & -d_i(t) + h_d + c_0 f(d_i(t)) \\ & - c_1 \sum_{i' \neq i} f(d_{i'}(t)) + c_2 f(d_{i-1}^m(t)) \\ & - c_3 f(d_i^m(t)) - I_C(t) \end{aligned} \quad (4)$$

$$\begin{aligned} \tau \dot{d}_i^m(t) = & -d_i^m(t) + h_m + c_4 f(d_i^m(t)) \\ & - c_5 \sum_{i' \neq i} f(d_{i'}(t)) + c_6 f(d_i(t)). \end{aligned} \quad (5)$$

To understand the role of different terms here, it is useful to examine the state of this system at three critical moments in time during a transition illustrated in Fig. 2 as well as the associated time courses and dynamics shown in Fig. 3. The first three terms in Eq. (4) are a dynamic stabilization factor $-d_i(t)$, a negative resting level, $h_d < 0$, and a self-excitatory term of strength, c_0 . Herein, $f(\cdot)$ is the sigmoidal non-linearity of Eq. (3). Together, these terms establish a bistable dynamics of each ordinal node which can be either in an “off” attractor ($d_i < 0$) or an “on” attractor ($d_i > 0$) (see the top row of Fig. 3). Inhibition among all ordinal nodes of strength, c_1 , allows only one node to be “on” at any time. The ordinal node at the i th step is excited with strength, c_2 , when the memory node of the previous, $(i - 1)$ th, step is “on”, but is inhibited with strength, c_3 , when the memory node at the same, i th, step is “on”. The first of these inputs drives activation of the next step, while the second input helps to destabilize the predecessor of the new action. Loss of stability of the activated predecessor state requires in addition an inhibitory input, $I_C(t) = c_{CoS} \cdot \int_y f(U_C(y, t)) dy$, from the condition of satisfaction field, $U_C(y, t)$, Eq. (8).

Memory nodes have the same basic bi-stability between “on” and “off” states established by the first three terms of Eq. (5). The memory nodes are inhibited with strength, c_5 , by any activity in the ordinal layer, except for activity at the corresponding ordinal node, which excites with strength, c_6 . The strength, c_4 , of self-excitation is chosen such that the “on” state is sustained after this excitatory input has been removed, so that the memory node stays “on” after the sequence has advanced to the next step.

4.2. Action fields and synaptic projections

Ordinal nodes specify an action by projecting onto an action field, a dynamic representation of a continuum of possible actions spanned by an appropriate continuous parameter, x . The dynamics of the action field, Eq. (6), is set up by the first three terms, which define the generic neural field dynamics of Eq. (1).

$$\begin{aligned} \tau \dot{U}_A(x, t) = & -U_A(x, t) + h_A + \int f(U_A(x', t)) w(x - x') dx' \\ & + c_{ord} \sum_{i=0}^{N_d} f(d_i(t)) M_i'(x, t) + c_{vis} I_{vis}(x, t). \end{aligned} \quad (6)$$

The projection from the ordinal nodes, $f(d_i(t))$, is mediated by connection weights, $M_i'(x, t) = M_i(x, t) + \mathbf{1}$, where the last term provides for a global boost of the action field by any activity in the ordinal set. $M_i(x, t)$ is a matrix of modifiable connection weights that specify the connection strength between the i th ordinal node and the site, x , of the action field. The sum is over all N_d ordinal nodes, c_{ord} is a constant controlling strength of the ordinal contribution. The last term is a perceptual input whose strength is controlled by the factor, c_{vis} .

If the non-zero connection weights, $M_i(x, t)$, are localized around a particular action characterized by the parameter, x_0 , then activation of the associated ordinal node ($f(d_i(t)) > 0$) pushes sites around x_0 above threshold and a self-stabilized peak is induced in the action field after a detection instability. That peak is a stable neural representation of the action linked to the i th step of the sequence. Projections from the action field onto the motor surface steer the system so that it executes an appropriate physical action. An exemplary implementation of such motor behavior is given in Section 5 “Robotic implementation”. Fig. 4 illustrates how the ordinal set determines the location of a peak in the action fields

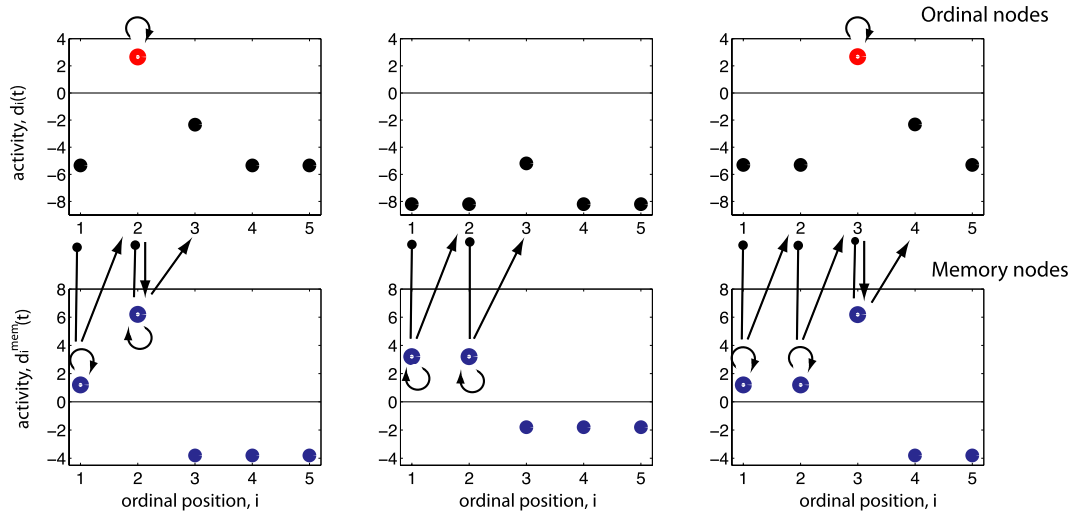


Fig. 2. The activation levels of the five ordinal (top) and memory (bottom) nodes at three time slices during a transition from the second to the third action in a sequence. *First column:* the state of the ordinal system while the second action in a sequence is produced, the second ordinal node is activated and specifies the ordinal position. *Second column:* the transition phase, when a condition of satisfaction signal inhibits the ordinal pool. The third ordinal node is less negative, because it receives input from the second memory node, but is not inhibited by the third memory node. *Third column:* after the condition of satisfaction signal ceases, the least negative third ordinal node reaches the threshold for a detection instability and gets activated. Arrows show active excitatory connections, while lines ending with a filled circle show active inhibitory connections.

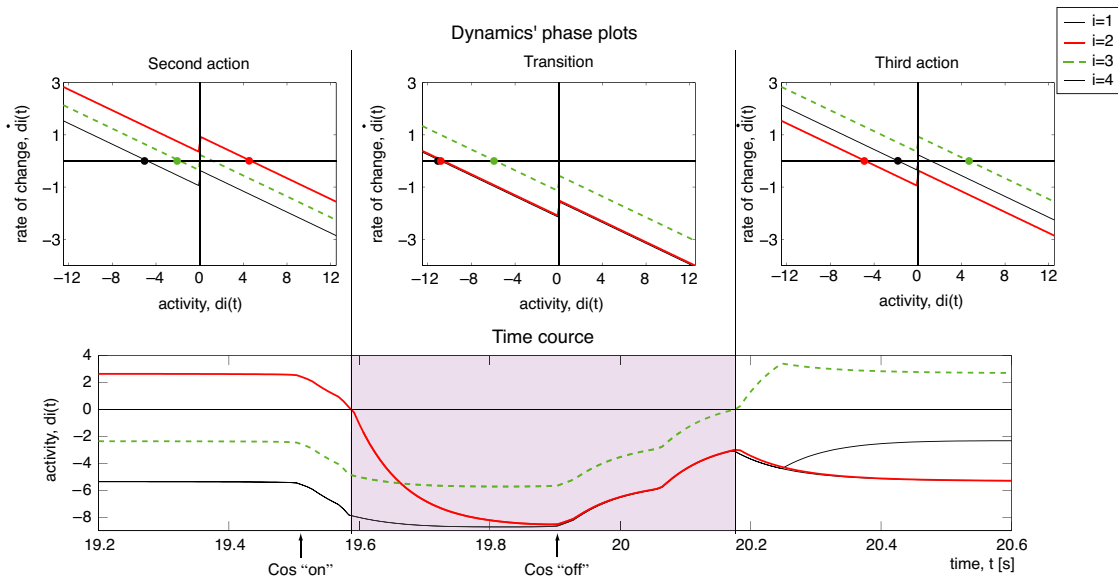


Fig. 3. *Top:* Visualization of the ordinal dynamics at the same three time slices during a transition, as in the Fig. 2. Intersections with abscissa, marked by circles, are stable fixed points corresponding to “off” (negative half-space), or “on” (positive half-space) states. *Bottom:* The time course of ordinal nodes’ activations during a sequential transition. The points in time when the condition of satisfaction signal enters the ordinal set (CoS “on”) and ceases (CoS “off”) are marked by arrows. The shaded region corresponds to the transition phase between the two actions.

by comparing two moments in time corresponding to the second (left) and the third (right) steps.

The weights of the neural connections, $M_i(x, t)$, are modified during learning according to a Hebbian-like rule, Eq. (7):

$$\tau_i \dot{M}_i(x, t) = (-M_i(x, t) + f(U_A(x, t))) \cdot f(d_i(t)). \quad (7)$$

Thus, when an ordinal node is active ($f(d_i(t)) > 0$), the weights are strengthened to those sites of the action field at which positive activation ($f(U_A(x, t)) > 0$) represents the currently demonstrated action. Below we will show how sensory input during learning induces localized activation peaks in the action field. Converging to the output of the action field on the timescale of learning dynamics, τ_i , the weight matrix, $M_i(x, t)$, assumes this localized form, so that a single action is associated with each ordinal node.

4.3. Condition of satisfaction field

For each action in a sequence, its termination criterion, or *condition of satisfaction* (CoS), is represented in a dynamic field, $U_C(y, t)$, defined over a dimension, y , that captures the sensory states identifying the terminal criterion of an action and need not be identical to the characteristic dimension, x , of the action field, Eq. (6). The CoS neural field evolves according to a dynamical equation Eq. (8):

$$\tau \dot{U}_C(y, t) = -U_C(y, t) + h_C + \int f(U_C(y', t)) w(y - y') dy' + c_{AT}(x, y) * f(U_A(x, t)) + c_{vis}^C I_{vis}(y, t). \quad (8)$$

The first three terms here set up the usual Amari field dynamics with time constant, τ , resting level, h_C , and interaction kernel,

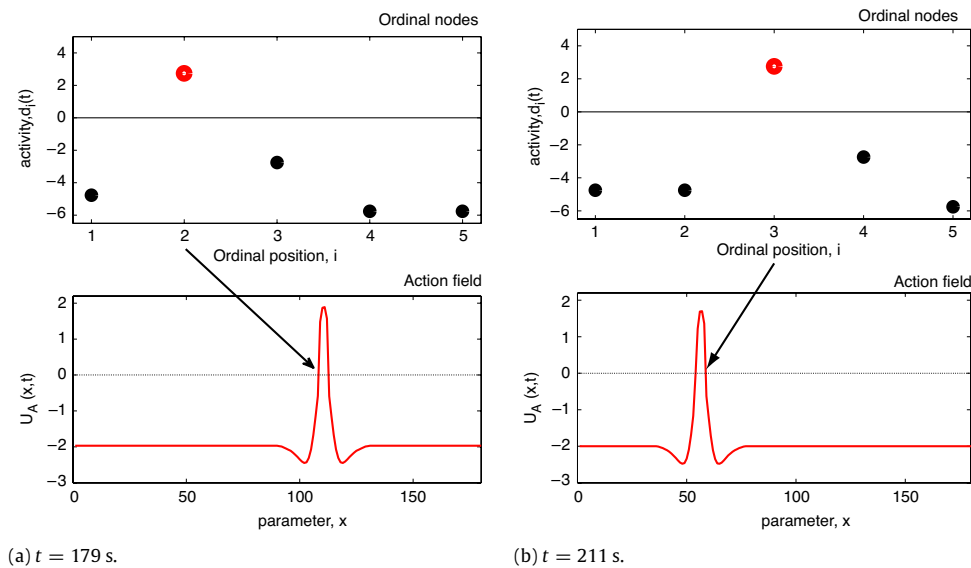


Fig. 4. Projections from the ordinal nodes onto the dynamic field representing a neural parameter of action. (a) At time $t = 179$ s, the second ordinal node, $d_2(t)$, is active and induces the representation of the associated action in the action field, $U_A(x, t)$. (b) At time $t = 211$ s, the third ordinal node is active and induces a different peak in the action field that represents the third action. Arrows indicate centers of the learned localized projections from the ordinal nodes to the action field.

$w(y - y')$. Positive activation in the action field ($f(U_A(x, t)) > 0$) propagates to the CoS field through the synaptic mapping $T(x, y)$ which defines the mapping between the dimensions characterizing actions, x , and their terminal states, y . Sensory input, $I_{vis}(y, t)$, of strength, c_{vis}^C , induces peaks in the CoS field when it overlaps in the field dimension, y , with input from the action field; c_A is a constant.

The function of the CoS field is illustrated in Fig. 5. A peak of activation in the action field (panel A) represents a particular action and provides localized input to the CoS field (B), making it sensitive to sensory input (C) at matching sites. When the terminal state of the action is detected by the sensory system (F), sensory input to the CoS field overlaps with the preshaping input from the action field (D). The integrated inputs are sufficient to surpass the activation threshold of the CoS field and a self-stabilized peak arises (E). The positive activation in the CoS field signals the successful accomplishment of the on-going action. This neural signal is propagated to the ordinal set, where it uniformly inhibits all ordinal nodes and suppresses any positive activity in the ordinal systems. This triggers an instability in the dynamics of the system, which is described in the following subsection.

4.4. The sequential transition as a cascade of instabilities

The concurrent evolution of the ordinal system, the action field, and the CoS field linked to the environment through motor behavior and sensory signals leads to the sequential transition from one action to the next in a cascade of instabilities in the dynamical sub-systems of the architecture that are illustrated in Fig. 6.

We look at the system when the second action in the sequence has been activated and controls the motor behavior. The cascade of instabilities begins when the terminal condition of that action is picked up by the sensors. This leads to positive activation in the perceptual system, which feeds into the CoS field. There, the perceptual input matches the preactivation from the action field. The combined input is capable of pushing the CoS field through the detection instability, leading to a self-stabilized peak. The new CoS peak inhibits the ordinal system, pushing all nodes including the activated second one below threshold (a reverse detection instability). As a result, input is removed from the peak in the

action field, pushing it beyond the reverse detection instability, so that this peak decays. Its decay removes critical input from the CoS field, triggering a reverse detection instability there, so that the CoS peak decays. This removes global inhibition from the ordinal system. Within the ordinal system, the third node has a competitive advantage because it is boosted by the memory node of the second step. It goes through the threshold first, stabilizes its “on” state, and inhibits the other ordinal nodes. The projection of the third node onto the action field pushes that field through the detection instability at the location associated with the third node. The new peak in the action field impacts on the motor behavior of the system. The system has the equivalent dynamic state as at the beginning of the cascade, but now with the third rather than the second sequential action in execution.

Note that each instability within a cascade takes place within one of the dynamic fields or discrete neural dynamics. In terms of the complete overall dynamics, a cascade is a single transient triggered by the detection instability in the CoS field. The robust temporal order of the instabilities within the cascade is established by the structure of the neuronal architecture.

The same cascade of instabilities is at work during sequence learning. During learning, the peak in the action field is generated by direct sensory input that originates from a demonstration of the action. The co-activation of an ordinal node and the action field drives the synaptic weights from the ordinal node to the action field. Learning of a step in the sequence is terminated by generating a sensory signal that feeds into the CoS field. This sensory signal can be delivered by an autonomous action segmentation mechanism. This triggers the same cascade of transitions illustrated above and makes the system ready to learn the next sequential action. An exemplary implementation of the motor and perceptual systems, and demonstration of components' integrated dynamics guiding a real-world action are presented next.

5. Robotic implementation

To demonstrate how the DFT sequence generation architecture can be linked to simple sensory and motor systems and thus be embodied, we implement the architecture on an autonomous robot in a simple task setting, a sequential color-search task. Presenting differently colored objects to the robot's camera in a given order, the

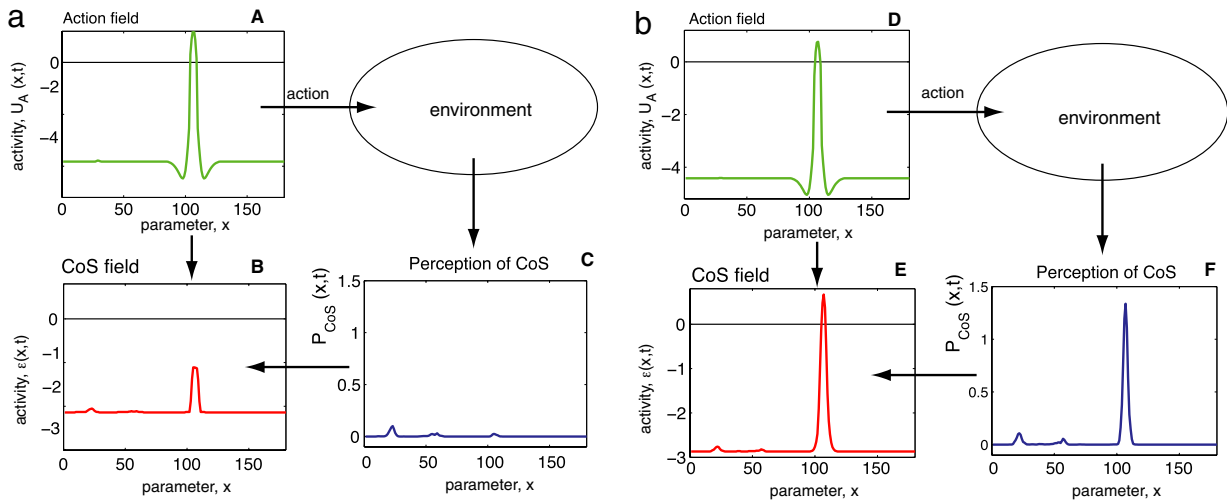


Fig. 5. Detection instability in the condition of satisfaction (CoS) field. (a) During the action production, the CoS field (B) is preshaped by the action field (A) to be sensitive to the desired end-state of an action, that can be detected by the perceptual module (C). (b) Beginning of the transition phase: When sensory input (F) matches the preshape from the action field (D), a peak in the CoS field (E) is induced and stabilized.

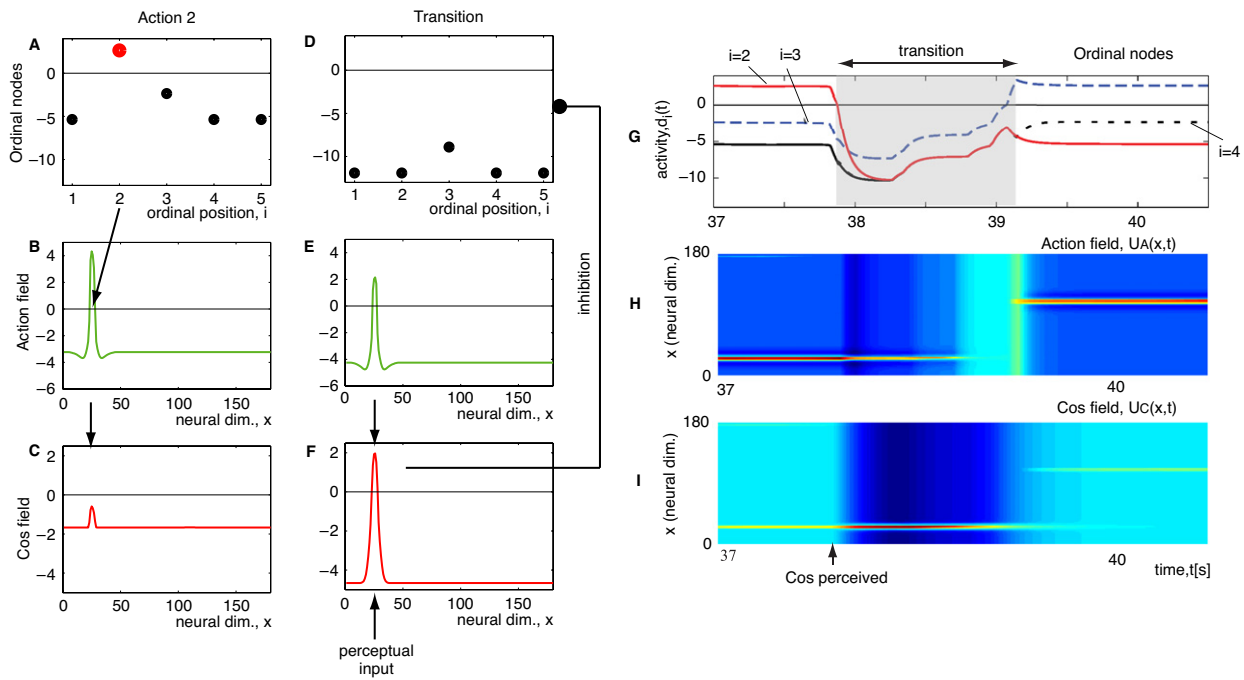


Fig. 6. A transition between two actions in a sequence. *Left:* The activation of the ordinal nodes, the action field, and the condition of satisfaction (CoS) field during the second action. The second ordinal node (A) projects onto a location, $x \approx 30$, of the action field (B), and induces a self-stabilized peak there that represents the on-going action. The CoS field (C) is pre-activated at the matching location. *Middle:* The same variables after a sensory signal has been perceived that matches the terminal condition of the second action. This induces a peak of activation in the CoS field (F), which inhibits the ordinal pool (D). The peak in the action field decays (E). *Right:* Time courses of the ordinal nodes (G), the action field (H), and the CoS field (I). The two fields are represented through a color code, in that red indexes positive and blue strongly negative levels of activation, light blue indexes activation close to the resting levels of the fields, yellow indexes low levels of positive activation. The first event is the creation of a new peak in the CoS field (detection instability). This triggers the inhibition of all ordinal nodes below threshold and leads to the decay of the peak in the action field (reverse detection instability). Next, the peak in the CoS field decays (reverse detection instability), which leads to the release of the ordinal system from inhibition and the activation of the third ordinal node. Finally, a new peak is formed in the action field (detection instability) that represents the third action. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

system is “taught” a sequence of colors (Fig. 7(a)). Here, learning is not supervised in the traditional sense in that the sensory signal that drives learning is not an error signal that informs about what would have been the correct response. Instead, the user’s supervision consists of demonstrating the sequence by creating environmental conditions conducive to generating relevant sensory inputs. The acquisition of these inputs is autonomous.

Controlled by a simple behavioral dynamics that comprises target acquisition and obstacle avoidance, the robot is capable of searching for an object of a given color within an arena, in which

colored objects are distributed (Fig. 7(b)). The robot does so in the order in which the colors were presented during learning. For instance, if taught the sequence “red–blue–green–red–yellow”, the robot will first search for a red object, then for a blue object, then for a green object, and so on. Each time an object of the currently requested color looms large on the robot’s visual array, that object is considered “found” and the robot switches to the next color in the sequence.

This simple random search scenario demonstrates core properties of the DFT sequence generation architecture: (1) the capacity




Fig. 7. Robotic implementation of the sequence generation architecture on a Khepera robot vehicle equipped with a color vision system. (a) A sequence of colors is “taught” to the robot by presenting colored objects in a particular order to the robot’s camera. (b) The robot is then capable of navigating the arena, locating and approaching colored objects in the learned order.

to derive and maintain stable representations of action goals from a simple sensory system, (2) the capacity to control real-world motor behavior, and (3) the capacity to obtain a reliable sensory signal that controls the switch from an action to its successor. Because it takes variable amounts of time to find an object of the requested color, this task highlights (4) the capacity of the sequence generation system to stabilize an action goal until it has been achieved. Finally, the learning phase demonstrates (5) the capacity to autonomously acquire the serial order of a sequence from sensory signals.

5.1. The action and condition of satisfaction systems

The action and condition of satisfaction fields work on the basis of fairly low level sensory information. Because color is the perceptual dimension that specifies which action within a sequence is currently activated, the action field is defined over hue as the dimension, x . A supra-threshold peak of activation in the action field provides ridge-shaped input to a perceptual color-space field defined below. Such input is localized along the color dimension, but constant along space.

The condition of satisfaction field should generate a peak when a block is approached whose color matches the hue values currently activated in the action field. The condition of satisfaction field is, therefore, defined over the same hue dimension, labelled y when referring to the condition of satisfaction field. Formally, the mapping, $T(x, y)$, in Eq. (8) is the identity in this simple example. A peak in the action field centered over the hue value, x_0 , provides localized input to the condition of satisfaction field around that same value. Perceptual input to the condition of satisfaction field is received from a central portion of the camera image. The dynamics of this field is tuned such that the detection instability is reached only when a sufficiently large blob of pixels register hue values that match the input from the action field.

5.2. The color-space field: interface to sensors and motors

The two-dimensional color-space field (Fig. 8(A)) plays a dual role. On input, this field performs a parallel visual search, generating a peak at a spatial location at which the hue of a contiguous blob of pixels in the camera image matches the localized in hue input from the action field. On output, this field provides the target direction for the robot’s movement towards the location at which the color match has been detected. The color-space field is defined over dimensions of hue and horizontal axis of the image plane (the horizontal spatial dimension is sufficient to control the vehicle’s

heading direction). Fig. 8 shows how the sensory input to the color-space field is computed from the camera image: The distribution of hue values within a vertical column in the image, Fig. 8(B), is computed at each location, x_{im} , along the horizontal axis of the image plane, Fig. 8(C).

During sequence production, the color-space field is tuned such that a peak can be generated only when input from the camera overlaps sufficiently with ridge input from the action field that represents the color currently searched for (see the ridge along the spatial dimension in Fig. 8(A)). Sufficiently strong and spatially focused input from the camera, that falls onto the ridge, drives the field through the detection instability. The emerging peak signals that a suitably colored object has been found within the visual array. This peak then sets an attractor for a dynamics of heading direction that controls the robot’s movement (see Appendix A for details). As a result, the robot turns to and moves towards the suitably colored object. As the robot approaches the object, its visual projection onto the image plane grows in size until it triggers a peak in the condition of satisfaction field and the transition to the next color task in the sequence unfolds.

The strength of attraction toward the target direction scales with the size of the peak in the color-space field. The dynamics of heading direction also receives contributions from distance sensors for obstacle avoidance. In the absence of a peak in the color space field, these obstacle avoidance contributions are thus alone in determining the robot’s movement. In this mode, the robot will effectively wander around within an enclosed arena, generating what amounts to random search behavior. The velocity of the robot is controlled by a separate dynamics, that depends on the measured distance values, slowing the robot down in the vicinity of obstacles.

During sequence learning, a homogeneous input (boost) to the color-space field ensures that the input from the camera alone is sufficient to induce a peak in the color-space field. Such a peak represents the detection and, in the presence of multiple colored objects, selection of a colored object. The color information is passed on to the action field. Specifically, for each hue value, x , the maximal value of activation in the color-space field along the spatial dimension, x_{sp} , is fed into the action field. In the learning mode, the central portion of the camera image also provides input to the condition of satisfaction field. When a colored block is shown to the robot sufficiently close and centered on the robot’s camera, the condition of satisfaction field generates a peak that inhibits the ordinal set and thus stops the strengthening of the synaptic weights. Due to the high resting level of the perceptual field in the learning mode, the peak in the CoS field remains stable until visual input ceases, for instance, because the teacher removes the colored

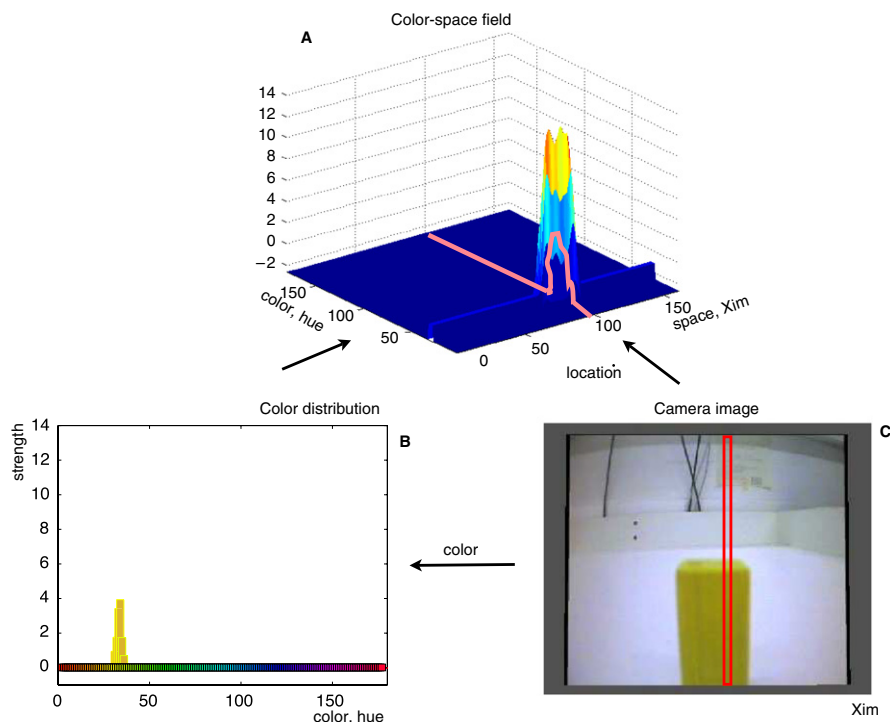


Fig. 8. The perception of the robot: for each column in the camera image (C), the color distribution (B) is determined. This distribution forms input to the location in the perceptual color-space field (A) that corresponds to the position of the column in the image. A large enough blob contiguous in both color and space induces a peak of activation in the color-space field, which is the stabilized representation of the sought object in the image-based reference frame.

block from the field of view of the robot's camera. The transient to the next item in the sequence is then accomplished. When a new object is presented, its color representation is associated with the next ordinal node.

6. Results: the model in action

6.1. Sequence learning

Fig. 9 illustrates how the robot is taught a sequence of colors. When the first block is presented to the robot's camera (Fig. 9, first column), the prevalent color in the image induces an activation peak in the perceptual color-space field. This peak is projected onto the color dimension of the action field and induces a peak there. Note that this reversal of the direction of coupling compared to the sequence production mode is brought about merely by the boost of activation to the perceptual color-space field. The presence of a peak in the action field, while a particular ordinal node is active, strengthens the synaptic connections from that node to the activated sites of the action field.

The activation peak in the action field also provides localized input to the condition of satisfaction field. Therefore, as soon as perceptual input to the condition of satisfaction field is sufficiently strong – because the object is brought sufficiently close to the camera – an activation peak is induced in the condition of satisfaction field. This peak then inhibits the ordinal set (Fig. 9, second column), leading to activation in the ordinal system to drop below threshold. The strengthening of synaptic weights stops at this point. When the teacher removes the colored block from the robot's view (Fig. 9, third column), the activation peaks in the action and the condition of satisfaction fields decay completely, releasing the ordinal set from inhibition. Due to the dynamics of the ordinal and memory nodes, the next ordinal node becomes activated.

When the next block with the next color is presented to the robot, the new color is detected in the color-space field

(Fig. 9, last column) and passed on to the action field in the same fashion. Synaptic connections between the next ordinal node and the activated sites of the action field are strengthened. The procedure continues until all color blocks have been presented to the robot. The outcome of the learning procedure is a set of learned connection weights from the ordinal nodes to the action field, which may now direct the flow of activation during sequence generation.

6.2. Sequence production

Fig. 10 illustrates how sequential color search behavior is produced by the robot. An active ordinal node induces an activation peak in the action field through the learned synaptic connections (Fig. 10(A)). The location of this peak specifies the color that must now be searched (here, "green"). The localized output of the action field sends ridge-shaped input to the perceptual color-space field (Fig. 10(B)), facilitating activation of that field at "green" sites. Thus, the locations of green objects in the camera image (Fig. 10(C)) compete for activation in the color-space field. A location receiving input from the largest green object surpasses the activation threshold first. Lateral inhibitory interaction within the color-space field selects and stabilizes the representation of this object (Fig. 10(B)). This activation peak in the color-space field now controls the robot's movement by specifying an attractor in its heading direction dynamics. As the robot moves, the shifting location on the sensory surface of the visual projection of the selected green object is tracked by the peak in the color-space field.

The activation peak in the action field pre-activates the condition of satisfaction field (Fig. 10(D)), making it sensitive to "green". When the robot has approached the target block, the associated color blob takes up a large portion of the image (Fig. 10(H)). As a result, perceptual input to the condition of satisfaction field, summed with the prior input from the action field, surpasses the activation threshold of the detection instability (Fig. 10(I)). The emerging activation peak in the condition of satisfaction field signals the

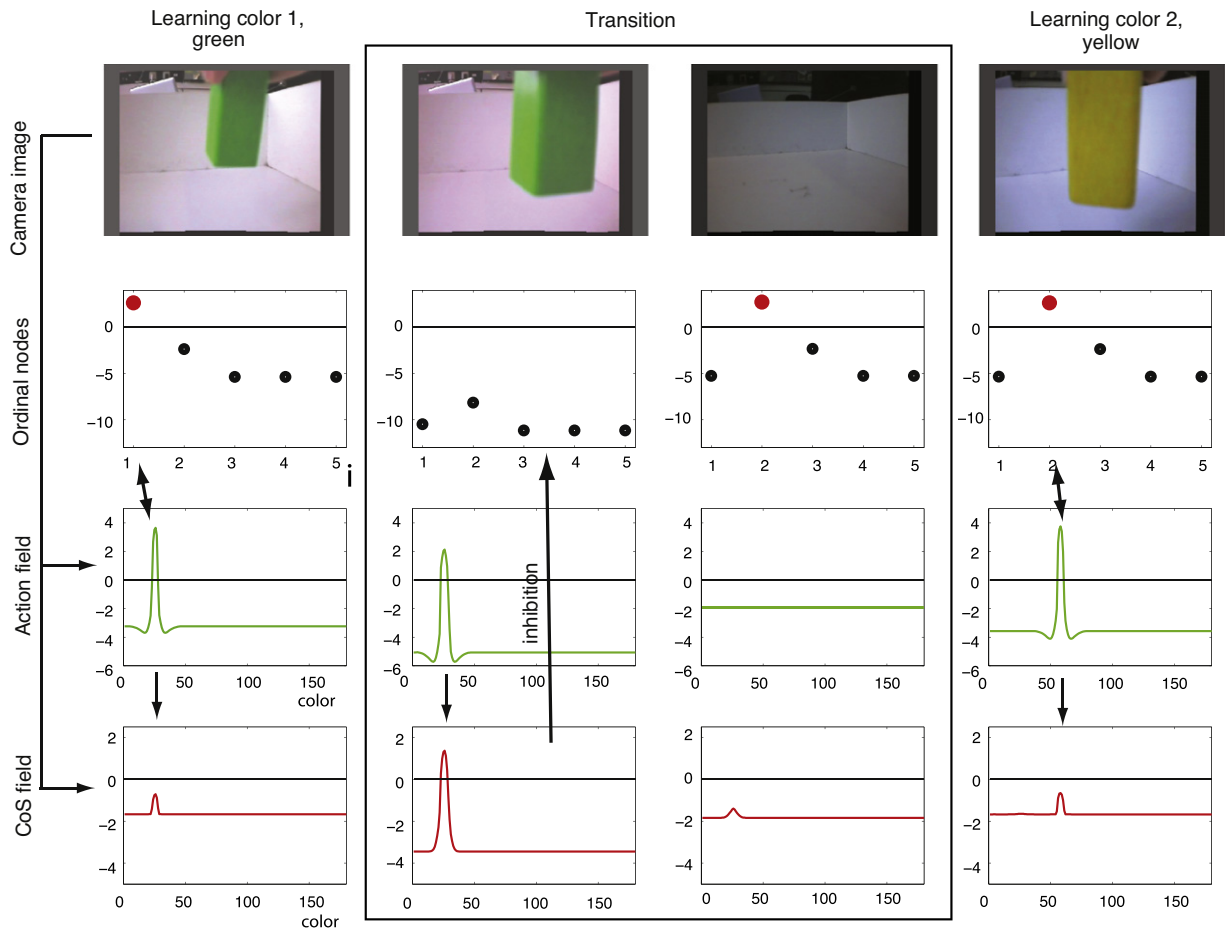


Fig. 9. Sequence learning in the DFT architecture, driven by a user-presented visual input. See text for details.

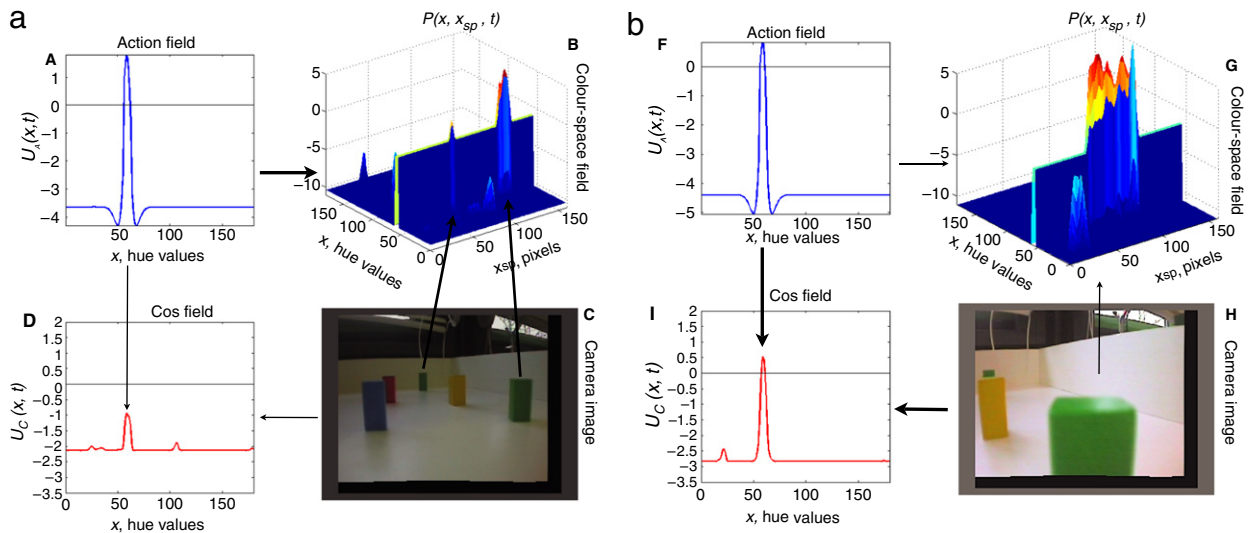


Fig. 10. Two snapshots of the dynamics of the DFT sequencing architecture during sequence production on a real robot. See text for details.

successful accomplishment of the color-search action at this stage of the sequence.

The peak in the condition of satisfaction field inhibits the ordinal nodes, triggering the cascade of transition instabilities. The peak representing the current action in the action field decays, which causes decay of the peak in the condition of satisfaction field. Ultimately, the representation of the next action in both ordinal and action systems is activated. The previous color loses its

advantage in the color-space field, where the peak decays. Consequently, the attractive force in the heading direction dynamics wanes. The green block now only acts as an obstacle, helping the robot to get on its way towards searching for the next color.

Fig. 11 shows the time course of the dynamics of the ordinal nodes, the action field, the condition of satisfaction field, and the spatial projection of the color-space field during learning and production of a sequence “red–blue–green–blue–yellow” (RBGBY).

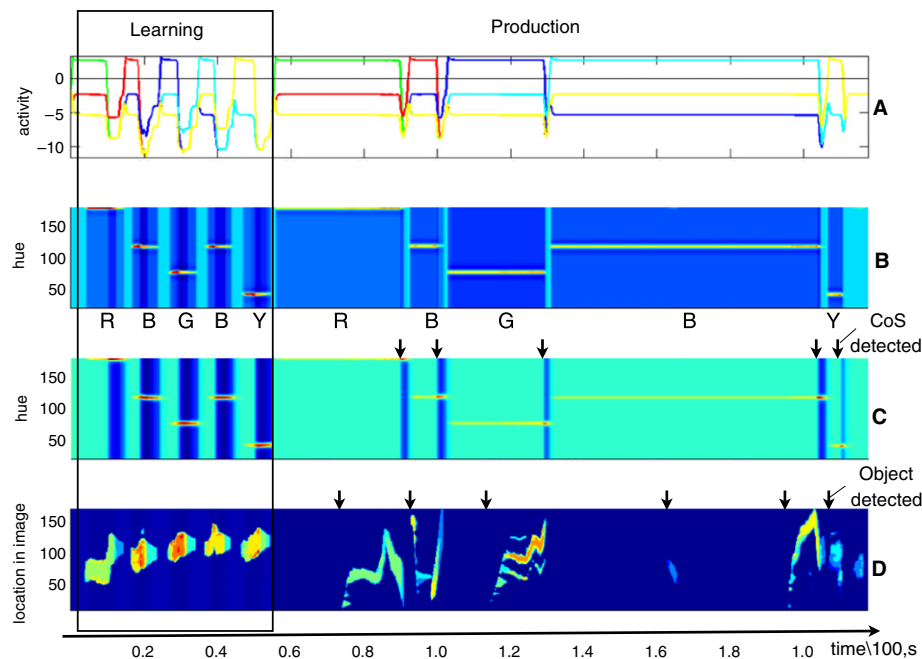


Fig. 11. One run of the robotic demonstrations. (A): Time courses of activation of five ordinal nodes during sequence learning and production. (B): Time course of activation in the action field. Positive activation in the field encodes the color currently searched for. (C): Time course of activation in the condition of satisfaction field. Arrows mark the times when condition of satisfaction signals were emitted (detection instabilities in the field). (D): The projection of the perceptual color-space field onto the spatial dimension (horizontal axis of the image plane). The arrows mark times when the object of interest in each ordinal position first appeared in the visual array of the robot. The “random search” behavior changed to “approach target” behavior at these points. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In the following, we present several runs of the robot system in different environments, merely varying which sequence of colors is taught during sequence learning and the physical arrangement of objects in the arena during sequence production. Exactly the same parameter setting is used for the neural dynamics in all demonstrations.

6.3. Timing of actions

The core property of the DFT architecture is the stability of action representations at each ordinal position enabling the system to tolerate variable durations of the individual actions when the sequence is executed in an unknown environment. In the simulations illustrated in Fig. 12, the durations of actions range from a few seconds to two minutes due to different spatial arrangements of obstacles and targets.

6.4. Flexibility of sequence generation: no problems with repetitions

The DFT architecture does not make use of direct connections between the representations of actions at different stages of the sequence (no “chaining” mechanism), nor is inhibition of the previous items essential for sequence learning and production. This makes it possible to learn any sequence, including those in which the same actions are repeated at different, even adjacent ordinal positions. This is demonstrated by the run illustrated in Fig. 12(a), in which the “blue” is searched at the second and fourth step. In the run illustrated in Fig. 12(b) the same color “green” is requested twice in a row. Disentangling ordinal information from the content of each action enables this form of flexibility.

6.5. Noisy environments

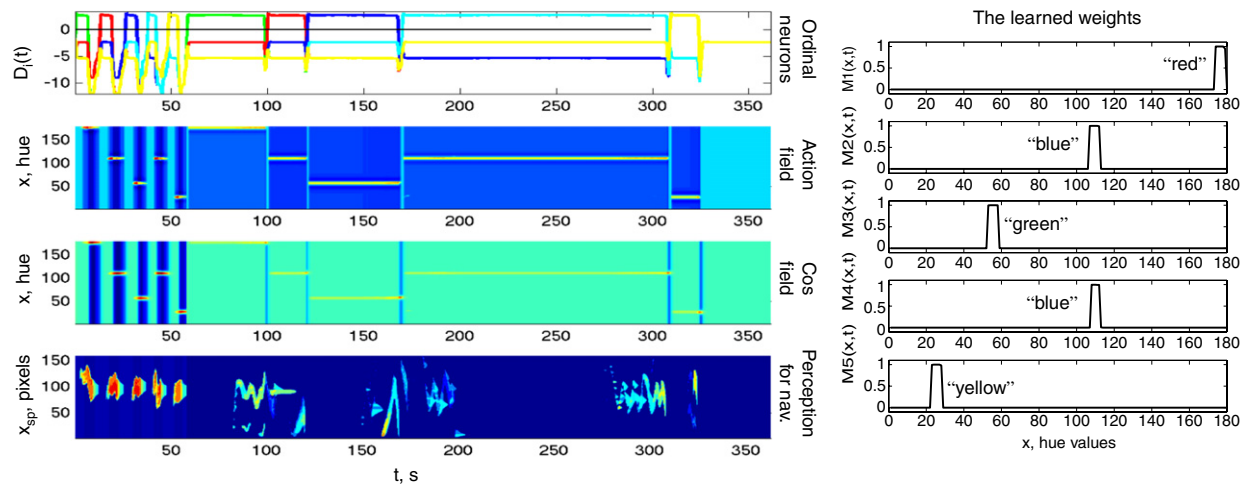
To illustrate the robustness of sequence generation in the DFT architecture, some sequences were acquired outside the arena. Thus, during the learning phase input was more complex, no longer dominated by a single prevalent color. Fig. 12(c) shows that the

system was still able to detect the most salient color during the learning phase and to learn the correct sequence. Note that a more complete architecture capable of representing objects and scenes (Faubel & Schöner, 2008) must replace the simple color-space perceptual field in order for this system to function robustly in complex, feature-rich environments.

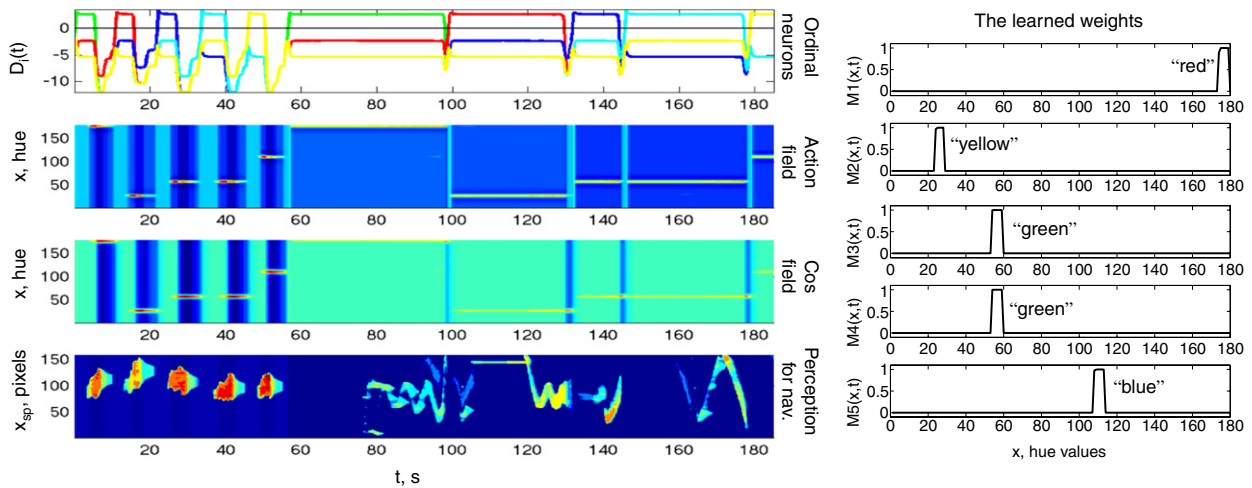
7. Discussion

How humans and cognitive robots generate sequences of actions is constrained by their embodiment and situatedness (Riegler, 2002). The physical properties of the body imply, for instance, that actions take characteristic amounts of time that may vary depending on circumstances. Physical environments may vary in time on their own characteristic timescale. Sensory and motor systems operate in continuous time and have graded state variables that are subject to spatio-temporal fluctuations. Sequence generation must be capable of integrating these constraints in order to produce stable behavior.

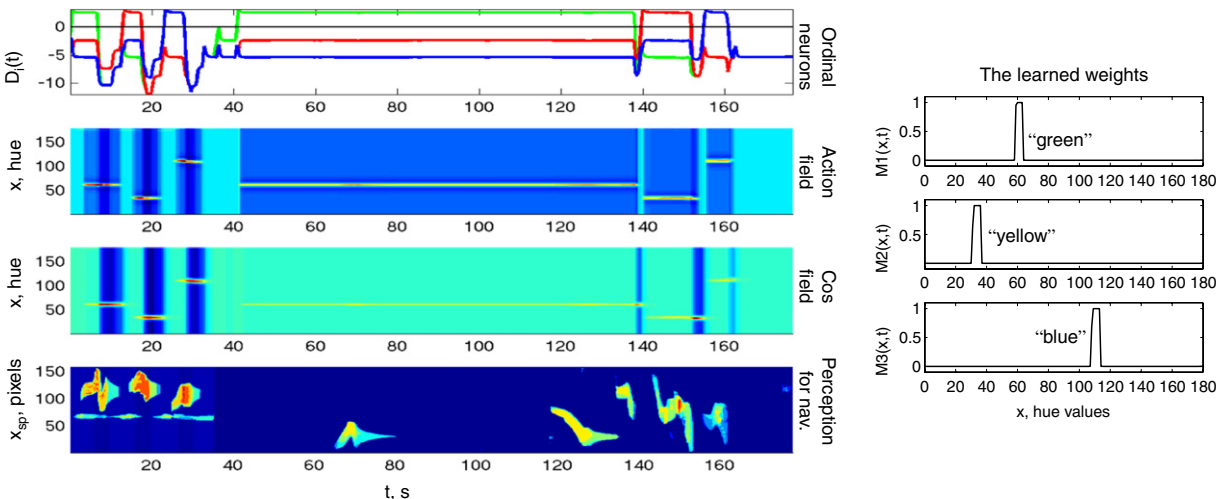
In this paper, we set out to develop a theoretical framework and model that address sequence generation under the constraints of embodiment and situatedness. Dynamic Field Theory (DFT) provides a theoretical language that captures principles of neural function, in particular, the autonomous, graded, and temporally continuous nature of neuronal processing. DFT resolves two inherently contradictory constraints for sequence generation that arise out of embodiment. On the one hand, any particular action within a sequence must be stable and resist change as environmental conditions vary. On the other hand, to bring about the transition to the next action, the current action must yield. In the DFT sequence generation model, action states are attractors that become unstable when sensory information signals successful completion of an action. In order for sensory information to play such a role, the switching signal itself must be represented stably. The condition of satisfaction system, inspired by Searle’s philosophical analysis of the structure of intentionality (Searle, 1983), provides such a stable



(a) Sequence RGBY acquired and produced in the arena.



(b) Sequence RGGY acquired and produced in the arena.



(c) Sequence GYB acquired outside and produced in the arena.

Fig. 12. Three demonstrations of sequence learning and production. *Left:* The time courses of activation in the functional modules of the DFT sequencing architecture. *Right:* The weights holding the sequence of colors after learning. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

representation. This likewise goes through instabilities, first when sensory information activates the condition of satisfaction signal, and then again when the successful transition to the next element of the sequence destabilizes the condition of satisfaction signal.

We have demonstrated the fluent execution of serially ordered actions in a physically embodied agent. The exemplary laboratory task included searching for a particular kind of object as a relevant action within each sequence. Searching exemplifies the problem of

varying time demands of each element of a sequence because the time needed to find a particular kind of object varies unpredictably. The capacity to stabilize a state of a sequence over variable durations also enabled our system to learn from demonstration without particular requirements for the timing of the teaching events. The teacher may present colored objects quite freely to the robot, the color of the object in view being detected and the connection weights being learned within a few milliseconds as determined by the time constant of the learning and field dynamics (see parameter values listed in Appendix B).

The concept of creating stable states at each stage of a sequence, which lose stability in a bifurcation, sets our approach apart from related neural dynamics models (Deco & Rolls, 2005; Rabinovich, Huerta, & Afraimovich, 2006; Selinger, Tsimring, & Rabinovich, 2003). These models share with our approach the commitment to principles of neural processing. They generate sequences of semi-stable states through carefully designed transient dynamics. The timing of the sequential transitions is controlled, however, by the internal dynamical properties of the neurons and thus resides at the neural timescale. Constraints that arise from acting out the behavioral sequences by embodied agents in time-varying environments are not addressed. This is true also for theoretical proposals for how to bridge, in principle, the gap between neural and behavioral timescales (Maass, Natschläger, & Markram, 2002; Melamed, Gerstner, Maass, Tsodyks, & Markram, 2004). The self-stabilization and switching mechanisms postulated for the encoding of serial order in a modular neural network model (Beiser & Houk, 1998) come closer in spirit to the DFT approach.

A conceptual approach to neural dynamics that is closely related to DFT is based on the notion of Hebbian cell assemblies. In that framework, ensembles of dynamical neurons are effectively bound together by strong excitatory interaction in the presence of global inhibitory coupling (Wennekers, 2006; Wennekers & Palm, 2007). Like DFT, this approach accounts for stable patterns of population activity. The concept of dynamic neural fields has been shown to emerge as a limit case when the coupling structure is homogeneous (Deco, Jirsa, Robinson, Breakspear, & Friston, 2008; Potthast & beim Graben, 2009). Within the framework of Hebbian cell assemblies, a model of sequence generation has been proposed that shares key ideas with our architecture (Wennekers & Palm, 2009). In that model, sequences may be generated transiently from semi-stable states of neuronal assemblies (“autonomous mode”), but may also emerge from actual attractor states. The attractors can be switched by a global signal, similar to the condition of satisfaction signal of our model.

Embodiment is only a secondary concern for the broader and more abstract neural dynamic models of sequence generation (Gnadt & Grossberg, 2007; Grossberg & Pearson, 2008) that are based on Grossberg's proposal of a neural mechanism for serial order (Grossberg, 1978). Here, sequences are stored as spatial gradients of neural activation. A volitionally-activated non-specific rehearsal wave controls the fluent unfolding in time of a sequence at production. An inhibition of return mechanism switches sequential elements off a fixed time after their activation. By varying the speed of movement generation through a graded “go” signal, the authors modeled variable timing of sequential arm movements including anticipatory preparation to the up-coming action (Grossberg & Pearson, 2008). Unexpected variation of the time needed to terminate an action was not addressed.

A series of neural dynamics models of motor control (Stringer, Rolls, & Taylor, 2007; Stringer, Rolls, & Trappenberg, 2004) have a considerable conceptual and mathematical overlap with our approach. The models are based on continuous attractor networks that could be viewed as particular mechanistic implementations of neural fields. Within this framework, associating subnetworks that represent a state with motor networks makes it possible to

learn arbitrary motor sequences (Stringer, Rolls, Trappenberg, & De Araujo, 2003). The sequences are performed autonomously, without feedback control, but at variable speeds producing different levels of force. To reach to the more flexible, cognitive level at which our approach is aimed, the motor control model would need to be embedded in an architecture that would include mechanisms for how information from the environment control an unfolding sequence.

The embodiment of sequence generation is, to a limited extent, addressed by architectures based on recurrent neural networks (Botvinick & Plaut, 2004; Dominey, Arbib, & Joseph, 1995; Elman, 1990). Here, the weights of a neural network are trained by exposing an agent to different behavioral sequences. The resulting network produces an appropriate sequence of outputs representing actions in a simulated environment. The environment provides input after each sequential action, which, combined with the internal state of the network, leads to the next output value. How that perceptual input is acquired from noisy environments is not addressed nor is the time course of action controlled and stabilized. Neural network models of this type have not been implemented in real-world agents and such implementation would require solving these problems. Relatedly, recurrent neural network architectures do not model the process of acquiring an action sequence in a natural setting with a single or few demonstrations. The back-propagation algorithm, that these networks rely on, requires extensive learning under supervision. The models do address, however, empirical findings on routine sequences and may provide insights into action slips (Botvinick & Plaut, 2004).

Naturally, the problem of embodiment is addressed within the domain of autonomous robotics and autonomous agents research. Here, sequence generation is more typically looked at as a problem in behavioral organization, in which rules determine the sequential ordering of actions (Steinhage & Schöner, 1998). From that perspective, serial order is the simplest case, in which any order is possible (although serial order includes the problem of learning such arbitrary sequential arrangements). A typical approach is to represent actions as discrete nodes at one or a number of hierarchical levels. Directed links between these nodes define the logical structure of behavioral organization, whereas inputs represent goals and environmental conditions. Such architectures have demonstrated behavioral sequencing in simulated environments and simple robotic demonstrations (Maes, 1989; Payton, Rosenblatt, & Keirse, 1990; Tyrrell, 1993). In order to apply these models in the robotic domain, the temporal continuity of the actions and their finite duration must be taken into account. One way to do that is to use event-driven processing in finite state machines which are in stationary states between transitions (Arkin & MacKenzie, 1994; Kosecka & Bajcsy, 1993). In this view, cognitive architectures are effectively decoupled from the sensory systems. The downside is that the stability problem is shifted to the level of sensory preprocessing, at which the relevant problems of sensor fusion, salient events detection, and segmentation of relevant objects are still broadly unsolved. Because many of these processes are actually cognitive in nature, their relegation to a system outside behavioral organization may be problematic.

The rich and complex flow of sequential actions in daily life is likely to be strongly constrained by the inherent logic of action (e.g., you cannot put down an object before grasping it), analogous to the rules of behavioral organization in robots. In cognitive psychology, however, sequence generation has more commonly been considered as a problem of serial order, that is, the arrangement of actions in an arbitrary but fixed order. The large literature on theoretical models of serial order aims at understanding the structure of the underlying memory systems and uses errors of serial order production as well as response times as constraints. These models are typically disembodied, that is, not concerned with the physical

production of each action and the acquisition of relevant sensory information. A useful classification of these models distinguishes between chaining, ordinal, and positional theories (Henson, 1998). Chaining theories postulate that serial order is stored in directional links between the successive states. A number of arguments can be advanced against chaining theories (Cooper & Shallice, 2000; Dell et al., 1997; Henson, 1998), although elements of chaining may be relevant to behavioral organization (Botvinick & Plaut, 2006; Hikosaka et al., 1999). Ordinal theories postulate that the order of an item in a sequence is represented in an activation gradient over a network that represents item information (Glasspool, Shallice, & Cipolotti, 1999; Grossberg, 1978; Page & Norris, 1998). Ordinal theories seem to contradict some experimental data on human (Henson, 1998) and animal (Eichenbaum, 2007) behavior, however. To date, positional models are the most successful candidates in addressing behavioral data on serial order (Burgess & Hitch, 1999; Dell et al., 1997; Henson, 1998; Houghton & Hartley, 1995). In positional theories, the item information is associated with a location that encodes ordinal position. This makes it possible to both account for ordinal errors through disruptions within the positional representation of serial order as well as to account for similarity effects and the influence of the domain structure.

The DFT architecture could be considered a neural implementation of a positional mechanism for encoding and producing serially ordered actions. Item information is represented by activation distributions within neural dynamical fields, whereas the serial order of items is represented separately by the location of the activated node within the ordinal network. Ordinal nodes project onto the neural fields, associating ordinal position with item information. This marks the contrast with chaining theories, in which item information of each step is directly linked to item information of the consecutive step.

Neurophysiology provides convergent evidence for such a positional account. Although much remains to be known about the neuronal basis of sequential behavior, neurons coding for ordinal position have been found in a number of relevant neuronal structures (Tanji, 2001). Neurons in the supplementary motor area were found to become active for specific movements within a movement sequence, but not when those movements were performed outside the sequence context (Shima & Tanji, 1998). In anterior cingulate cortex, some neurons showed activity selectively for a particular ordinal position, but independent of which particular movement was performed at the upcoming, preceding or subsequent stage of the sequence (Procyk et al., 2000). A similar pattern was found when investigating the difference in neural activations during “syntactic” and “natural” grooming sequences in rats (Aldridge & Berridge, 1998). A population of dorsolateral neurons in this study was active only during the particular sequential patterns of grooming movements, and not when the same patterns occurred outside the grooming sequence. These populations thus appeared to encode the serial order and not the motor properties of constituent movements. Neural pools responsive to serial information were even found in the motor cortex (Carpenter et al., 1999). The ordinal nodes of the DFT model are inspired by these neural signatures. Prefrontal regions play a role in the executive control of sequential action (Fujii & Graybiel, 2003). Prefrontal neural activity may sustain information about sequences between trials (Averbeck & Lee, 2007). Activity after a learning trial and before production of a sequence predicts which of a number of learned sequences is going to be initiated. Different patterns of neural activity were observed during the preparation of a movement sequence, before a particular ordinal position, and in specific time intervals between serial items (Tanji, 2001). Overall, these findings support the notion of a positional code for ordinal system.

The basal ganglia play a role in the temporal organization of actions. The basal ganglia neurons are active during both the

learned motor or cognitive tasks and during learning of novel motor behavior (Hikosaka, Takikawa, & Kawagoe, 2000). The inhibitory connections are established from the discrete regions of basal ganglia to separate cortical regions (Parent & Hazrati, 1995). As this structure is also considered to participate in reward anticipation mechanism, the function of the presented here condition of satisfaction system might resemble that of the basal ganglia structures.

Clearly, multiple mechanisms support the production of sequentially ordered actions in natural and artificial systems. The DFT architecture developed in this paper demonstrates how sequential neuronal states can be generated from a continuous-time dynamics that is coupled to graded, time-varying sensory inputs and controls the movement behavior of a physical agent situated in the real world. The robotic implementation forced us to take the perception–action loop seriously and thus to address the constraints that arise from embodiment and situatedness. This led to the concept of each stage of a sequence being instantiated as an attractor state that becomes unstable in a bifurcation under the control of a condition of satisfaction system. The switch to the next stage of the sequence thus occurs in a controlled fashion, enabling flexible timing. Internal simulation of the behavior that generates a condition of satisfaction makes it possible to emulate the fixed timing observed in a number of motor skills (Sandamirskaya & Schöner, 2006). Beyond serial order, we believe that the DFT model proposed here may provide the basis for a comprehensive account for behavioral organization and goal-oriented sequence generation.

Appendix A. Representation of the motor and perceptual systems

A.1. Motor system

In our example, the movement of the robot is controlled by dynamics of heading direction of the robot, Eq. (A.1) (Schöner, Dose, & Engels, 1995). The desired change of heading direction, $\Delta\theta = \dot{\theta}(t)\Delta t$, at each time step, t , of the dynamics' integration is translated into a difference of commands for the two wheels of the robot, and causes the desired rotation. The new heading direction, θ , is never computed explicitly here, but the robotic hardware, as part of the loop, “integrates” Eq. (A.1).

$$\tau_{\text{nav}}\dot{\theta}(t) = \lambda_{\text{obs}}F_{\text{obs},i}(t) + \lambda_{\text{tar}}F_{\text{tar}}(t). \quad (\text{A.1})$$

Here, λ_{obs} and λ_{tar} are constants controlling the strength of the target (attractor) and obstacles (repellers) contributions. The functions, $F_{\text{obs},i}$ and F_{tar} , describe the contributions to the dynamics from obstacles and the target perceived by the robot. The obstacle contribution is calculated with Eq. (A.2).

$$F_{\text{obs}} = \sum_i S_i(t) \cdot \exp\left[-\frac{\Psi_i}{2R_{\text{obs},i}(t)^2}\right]. \quad (\text{A.2})$$

Here, $S_i(t)$, $i = 1, \dots, 8$, are the signal strengths and Ψ_i are angular locations relative to the heading direction of the eight infrared sensors of the robot. $S_i(t)$ are calculated according to Eq. (A.3). $R_{\text{obs},i}(t) = \arctan\left[\tan\left(\frac{\Delta\Psi}{2}\right) + \frac{R}{R+D_i(t)}\right]$ is the angular range of an obstacle, that depends on the size of the robot, R , opening angle of the sensor, $\Delta\Psi$, and distance to the obstacle, $D_i(t)$. $D_i(t)$ is the distance to an obstacle in the direction of the i th sensor that is calculated from the sensor's value through an empirical calibration function, $D_i(t) = 0.9 \cdot \exp(-0.007 \cdot I_i(t))$ (mm), where β_1 , β_2 are numerical constants.

$$S_i(t) = \beta_1 \cdot \exp\left(-\frac{D_i(t)}{\beta_2}\right). \quad (\text{A.3})$$

The target contribution is calculated by extracting the spatial information from the output of the color-space field, $P(x, x_{\text{sp}}, t)$,

Table B.1

Numerical values of parameters used in implementation of the model. Units are displayed where applicable.

Name	Value	Name	Value
τ	10.0 (ms)	τ_{nav}	80.0 (ms)
τ_l	20.0 (ms)	τ_{vel}	30.0 (ms)
h_d	-5.0	λ_{tar}	0.2
h_m	-2.0	λ_{obs}	1.0
h_A	-1.8	β_1	15.0
h_C	-2.0	β_2	30.0
$c_{\text{inh}}^{A,C}$	5.0	Ψ_i	[-90, -45, -9, 11, 45, 90] (°)
β	100.0	v_{tar}	60.0 (pulse/s)
$c_{\text{exc}}^{A,C}$	0.9	v_{obs}	30.0 (pulse/s)
$\sigma^{A,C}$	3.0	R	100.0 (mm)
c_{CoS}	15	$\Delta\Psi$	45.0 (°)
c_{vis}	2.0	h_p	-2.0/ - 1.0 (learning)
$c_{\text{vis},C}$	1.0	c_A^P	2.5
c_{ord}	10.0	c_{vis}^P	2.0
c_0	7.2	c_{inh}^P	0.5
c_1	3.6	σ^P	[3.0, 3.0]
c_2	0.9	c_{exc}^P	0.6
c_3	0.8	c_p	1.0
c_4	3.5	c_A	1.5
c_5	2.0	N	180
c_6	2.6	N_{sp}	[160, 120]

Eq. (A.6). Thus, the target contribution is proportional to a distance from the peak of activity in this neural field to the center of the field, both calculated for the spatial projection, $P_{\text{pr}}(x_{\text{sp}}, t)$, of the two-dimensional neural field's output, Eq. (A.4).

$$F_{\text{tar}}(t) = \int_0^N P_{\text{pr}}(x_{\text{sp}}, t) \left(x_{\text{sp}} - \frac{N}{2}\right) dx_{\text{sp}}. \quad (\text{A.4})$$

Here, N is the (arbitrary) size of the perceptual color-space field, $P_{\text{pr}}(x_{\text{sp}}, t) = f(P(\arg \max_x (P(x, x_{\text{sp}}, t)), x_{\text{sp}}, t))$ is a projection of the output of the color-space field on the spatial dimension, x_{sp} ; x is the color dimension of the color-space field. Thus, integrating Eq. (A.1), the robot minimizes the angular distance between the current heading direction, θ , and the direction to the target, θ_{tar} , both of which, however, are not calculated explicitly.

The dynamics, Eq. (A.1), integrates impacts of the two sensor modalities—color vision and infrared sensing. The color vision is responsible for target acquisition and is in its turn affected by the dynamics of sequence generation. The sensori-motor dynamics, Eq. (A.1), thus provides for a low-level behavior organization.

The speed of the robot is controlled by Eq. (A.5):

$$\tau_{\text{vel}} \dot{v}(t) = -v(t) + (1 - \lambda_{\text{obs}}^v(t)) \cdot v_{\text{tar}} + \lambda_{\text{obs}}^v(t) \cdot v_{\text{obs}}, \quad (\text{A.5})$$

where $v_{\text{obs}} > v_{\text{tar}}$ are constant velocities with and without an obstacle in the vicinity of the robot respectively; λ_{obs}^v is a switch, that turns non-zero if an obstacle is perceived by an infrared sensor; τ_{vel} is the time constant of the dynamics.

A.2. Perceptual representation

The perceptual system in our example is the two-dimensional color-space field, Eq. (A.6), spanning the dimensions of color and horizontal position in the image. This field receives visual input from the camera of the agent: for each column in the image the color distribution is extracted and input to the corresponding row of the color-space field (see Fig. 8 and Section 5.2 “The color-space field: Interface to sensors and motors”).

$$\begin{aligned} \tau \dot{P}(x, x_{\text{sp}}, t) = & -P(x, x_{\text{sp}}, t) + h_p \\ & + \int f(P(x, x_{\text{sp}}, t)) w(x - x', x_{\text{sp}} - x'_{\text{sp}}) dx' dx'_{\text{sp}} \\ & + c_A^P f(U_A(x, t)) + c_{\text{vis}}^P V(x, x_{\text{sp}}, t), \end{aligned} \quad (\text{A.6})$$

where c_A^P is the constant characterizing strength of the input from the action field, $U_A(x, t)$, Eq. (6), that represents the ongoing action; $V(x, x_{\text{sp}}, t)$ is the visual input from the camera of the robot with amplitude c_{vis}^P . Other notations are the same as in Section 2 “DFT as a framework for embodied sequence generation”, the interaction kernel $w(x - x', x_{\text{sp}} - x'_{\text{sp}})$ is a two-dimensional symmetrical Gaussian.

A large continuous single-color blob in the image induces a localized peak of activation in the color-space field, that is the representation of a colored object at a certain position in the image. Location of the peak along the spatial dimension defines the contribution of the target for the navigation dynamics, Eq. (A.1).

Appendix B. Parameters used in implementation of the model

See Table B.1.

References

- Aldridge, J. W., & Berridge, K. C. (1998). Coding of serial order by neostriatal neurons: a “natural action” approach to movement sequence. *Journal of Neuroscience*, 18, 2777–2787.
- Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–87.
- Arkin, R., & MacKenzie, D. (1994). Temporal coordination of perceptual algorithms for mobile robot navigation. *IEEE Transactions on Robotics and Automation*, 10, 276–286.
- Averbeck, B. B., & Lee, D. (2007). Prefrontal neural correlates of memory for sequences. *Journal of Neuroscience*, 27, 2204–2211.
- Beiser, D. G., & Houk, J. C. (1998). Model of cortical-basal ganglionic processing: encoding the serial order of sensory events. *Journal of Neurophysiology*, 79, 3168–3188.
- Botvinick, M., & Plaut, D. C. (2004). Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychological Review*, 111, 395–429.
- Botvinick, M. M., & Plaut, D. C. (2006). Short-term memory for serial order: a recurrent neural network model. *Psychological Review*, 113, 201–233.
- Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, 107, 127–181.
- Burgess, N., & Hitch, G. J. (1999). Memory for serial order: a network model of the phonological loop and its timing. *Psychological Review*, 106, 551–581.
- Carpenter, A. F., Georgopoulos, A. P., & Pellizzer, G. (1999). Motor cortical encoding of serial order in a context-recall task. *Science (Reports)*, 283, 1752–1757.
- Clower, W. T., & Garrett, A. E. (1998). Movement sequence-related activity reflecting numerical order of components in supplementary and presupplementary motor areas. *Journal of Neurophysiology*, 80, 1562–1566.
- Cooper, R., & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, 17, 297–338.
- Deadwyler, S. A., & Hampson, R. E. (1995). Ensemble activity and behavior: what's the code? *Science*, 270, 1316–1318.
- Deco, G., Jirsa, V. K., Robinson, P. A., Breakspear, M., & Friston, K. (2008). The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Computational Biology*, 4, e1000092.
- Deco, G., & Rolls, E. T. (2005). Sequential memory: a putative neural and synaptic dynamical mechanism. *Journal of Cognitive Neuroscience*, 17, 294–307.
- Dell, G. S., Chang, F., & Griffin, Z. M. (1997). Connectionist models of language production: lexical access and grammatical encoding. *Cognitive Science: A Multidisciplinary Journal*, 23, 517–542.
- Dominey, P., Arbib, M., & Joseph, J.-P. (1995). A model of corticostriatal plasticity for learning oculomotor associations and sequences. *Journal of Cognitive Neuroscience*, 7, 311–336.
- Eichenbaum, H. (2007). Comparative cognition, hippocampal function, and recollection. *Comparative Cognition and Behavior Reviews*, 2, 47–66.
- Elman, J. (1990). Finding structure in time. *Cognitive Science*, 14, 179–211.
- Erickson, R. P. (1974). Parallel population neural coding in feature extraction. In F.O. Schmitt, & F.G. Worden (Eds.), *The neurosciences—third study program* (pp. 155–169).
- Ermentrout, B. (1998). Neural networks as spatio-temporal pattern-forming systems. *Reports on Progress in Physics*, 61, 353–430.
- Farrell, S., & Lewandowsky, S. (2002). An endogenous distributed model of ordering in serial recall. *Psychonomic Bulletin and Review*, 9, 59–79.
- Faubel, C., & Schöner, G. (2008). Learning to recognize objects on the fly: a neurally based dynamic field approach. *Neural Networks*, 21, 562–576.
- Fujii, N., & Graybiel, A. M. (2003). Representation of action sequence boundaries by macaque prefrontal cortical neurons. *Science*, 301, 1246–1249.
- Georgopoulos, A. P. (1991). Higher order motor control. *Annual Review of Neuroscience*, 14, 361–377.
- Glasspool, D. W., & Houghton, G. (2005). Serial order and consonant-vowel structure in a graphemic output buffer model. *Brain and Language*, 94, 304–330.

- Glasspool, D. D., Shallice, T., & Cicolotti, L. (1999). Neuropsychologically plausible sequence generation in a multi-layer network model of spelling. *Connectionist Models in Cognitive Neuroscience*, 40–51.
- Gnadt, W., & Grossberg, S. (2007). SOVEREIGN: an autonomous neural system for incrementally learning planned action sequences to navigate towards a rewarded goal. *Neural Networks*, 21, 699–758.
- Grossberg, S. (1978). A theory of human memory: self-organization and performance of sensory-motor codes, maps, and plans. In R. Rosen, & F. Snell (Eds.), *Progress in theoretical biology: Vol. 5* (pp. 233–374). New York: Academic Press.
- Grossberg, S., & Pearson, L. (2008). Laminar cortical dynamics of cognitive and motor working memory, sequence learning and performance: toward a unified theory of how the cerebral cortex works. *Psychological Review*, 115, 677–732.
- Hartley, T., & Houghton, G. (1996). A linguistically constrained model of short-term memory for nonwords. *Journal of Memory and Language*, 35, 1–31.
- Henson, R. N. (1998). Short-term memory for serial order: the start-end model. *Cognitive Psychology*, 36, 73–137.
- Hikosaka, O., Nakahara, H., Rand, M. K., Sakai, K., Lu, X., & Nakamura, K. (1999). Parallel neural networks for learning sequential procedures. *Trends in Neurosciences*, 22, 464–471.
- Hikosaka, O., Takikawa, Y., & Kawagoe, R. (2000). Role of the basal ganglia in the control of purposive saccadic eye movements. *Physiological Reviews*, 80, 953–978.
- Houghton, G. (1990). The problem of serial order: a neural network model of sequence learning and recall. In R. Dale, C. Mellish, & M. Zock (Eds.), *Current research in natural language generation* (pp. 287–319). London: Academic Press Professional, Inc.
- Houghton, G., & Hartley, T. (1995). Parallel models of serial behavior: Lashley revisited. *Psyche*, 2, 2–25.
- Johnson, J. S., Spencer, J. P., & Schöner, G. (2008). Moving to higher ground: the dynamic field theory and the dynamics of visual cognition. *New Ideas in Psychology*, 26, 227–251.
- Kosecka, J., & Bajcsy, R. (1993). Discrete event systems for autonomous mobile agents. *Robotics and Autonomous Systems*, 12, 187–198.
- Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–146). New York: Wiley.
- Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Computation*, 14, 2531–2560.
- Maes, P. (1989). How to do the right thing. *Connection Science Journal*, 1, 291–323.
- Melamed, O., Gerstner, W., Maass, W., Tsodyks, M., & Markram, H. (2004). Coding and learning of behavioral sequences. *Trends in Neurosciences*, 27, 11–14.
- Page, M., & Norris, D. (1998). The primacy model: a new model of immediate serial recall. *Psychological Review*, 105, 761–781.
- Parent, A., & Hazrati, L.-N. (1995). Functional anatomy of the basal ganglia. I. The cortico-basal ganglia-thalamo-cortical loop. *Brain Research Reviews*, 20, 91–127.
- Payton, D. W., Rosenblatt, J. K., & Keirse, D. M. (1990). Plan guided reaction. *IEEE Transactions on Systems, Man and Cybernetics*, 20, 1370–1382.
- Pfordresher, P., Palmer, C., & Jungers, M. (2007). Speed, accuracy, and serial order in sequence production. *Cognitive Science*, 31, 63–98.
- Potthast, R., & beim Graben, P. (2009). Dimensional reduction for the inverse problem of neural field theory. *Frontiers in Computational Neuroscience*.
- Procyk, E., Tanaka, Y. L., & Joseph, J. P. (2000). Anterior cingulate activity during routine and non-routine sequential behaviors in macaques. *Nature Neuroscience*, 3, 502–508.
- Rabinovich, M., Huerta, R., & Afraimovich, V. (2006). Dynamics of sequential decision making. *Physical Review Letters*, 97, 188103.
- Riegler, A. (2002). When is a cognitive system embodied? *Cognitive Systems Research*, 3, 339–348.
- Sandamirskaya, Y., & Schöner, G. (2006). Dynamic field theory and embodied communication. In I. Wachsmuth, & G. Knoblich (Eds.), *Lecture notes in artificial intelligence: Vol. 4930. Modeling communication with robots and virtual humans* (pp. 260–278). Springer.
- Schöner, G. (2008). Dynamical systems approaches to cognition. In R. Sun (Ed.), *Cambridge handbook of computational cognitive modeling* (pp. 101–126). Cambridge, UK: Cambridge University Press.
- Schöner, G., Dose, M., & Engels, C. (1995). Dynamics of behavior: theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems*, 16, 213–245.
- Searle, J. R. (1983). *Intentionality—an essay in the philosophy of mind*. Cambridge University Press.
- Selinger, P., Tsimring, L., & Rabinovich, M. (2003). Dynamics-based sequential memory: winnerless competition of patterns. *Physical Review E*, 67, 011905.
- Shima, K., & Tanji, J. (1998). Role for cingulate motor area cells in voluntary movement selection based on reward. *Science*, 282, 1335–1338.
- Steinhage, A., & Schöner, G. (1998). Dynamical systems for the behavioral organization of autonomous robot navigation. In G. T. McKee, & P. S. Schenker (Eds.), *Sensor fusion and decentralized control in robotic systems: proceedings of SPIE: Vol. 3523* (pp. 169–180). SPIE Publishing.
- Stringer, S. M., Rolls, E. T., & Taylor, P. (2007). Learning movement sequences with a delayed reward signal in a hierarchical model of motor function. *Neural Networks*, 20, 172–181.
- Stringer, S. M., Rolls, E. T., & Trappenberg, T. P. (2004). Self-organising continuous attractor networks with multiple activity packets, and the representation of space. *Neural Networks*, 17.
- Stringer, S., Rolls, E., Trappenberg, T., & De Araujo, I. (2003). Self-organizing continuous attractor networks and motor function. *Neural Networks*, 16, 161–182.
- Tanji, J. (2001). Sequential organization of multiple movements: involvement of cortical motor areas. *Annual Review of Neuroscience*, 24, 631–651.
- Tyrrell, T. (1993). The use of hierarchies for action selection. In J.-A. Meyer, H. L. Roitblat, & S. W. Wilson (Eds.), *From animals to animats: Vol. 2* (pp. 138–147). Cambridge MA, USA: MIT Press.
- Wennekers, T. (2002). Dynamic approximation of spatiotemporal receptive fields in nonlinear neural field models. *Neural Computation*, 14, 1801–1825.
- Wennekers, T. (2006). Operational cell assemblies as a paradigm for brain-inspired future computing architectures. *Neural Information Processing—Letters and Reviews*, 10, 135–145.
- Wennekers, T., & Palm, G. (2007). Modelling generic cognitive functions with operational Hebbian cell assemblies. *Neuronal Network Research Horizons*, 247.
- Wennekers, T., & Palm, G. (2009). Syntactic sequencing in Hebbian cell assemblies. *Cognitive Neurodynamics*, 3, 429–441.
- Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13, 55–80.

Yulia Sandamirskaya is a Ph.D. student in the Autonomous Robotics group at the Institut für Neuroinformatik at the Ruhr-Universität Bochum in Germany. Her research concerns development and robotic implementation of neural-dynamic models of embodied sequence generation and spatial language. She received a Diploma degree in Physics from the Belorussian State University (with honors) and a Master of Science degree from the Ruhr-Universität Bochum.

Gregor Schöner is the Director of the Institut für Neuroinformatik at the Ruhr-Universität Bochum in Germany and holds the chair for Theoretical Biology there. For the last 20 years, Gregor Schöner has brought his background in theoretical physics to bear on disciplines as varied as psychology, kinesiology, neuroscience, robotics, and computer vision. While working closely with experimental groups, he has developed theoretical concepts and models to understand movement, perception, and elementary forms of cognition on the basis of dynamical systems ideas. He has held academic positions in the US, France, and Germany, has published over 170 scientific articles and frequently lectures all over the world.