Dynamic Field Theory of Sequential Action: A Model and its Implementation on an Embodied Agent

Yulia Sandamirskaya, Gregor Schöner
Institut für Neuroinformatik
Ruhr-Universität Bochum, Germany

Published in *Proceedings of the IEEE 7th International Conference on Development and Leanring*, ICDL'2008, Monterey, California, USA (2008)

Dynamic Field Theory of Sequential Action: A Model and its Implementation on an Embodied Agent

Yulia Sandamirskaya
Institut für Neuroninformatik
Ruhr-Universität Bochum
44799 Bochum, Germany
Email: sandayci@neuroinformatik.rub.de

Gregor Schöner
Institut für Neuroninformatik
Ruhr-Universität Bochum
44799 Bochum, Germany
Email: gregor.schoener@neuroinformatik.rub.de

***Abstract*—How sequences of actions are learned, remembered, and generated is a core problem of cognition. Despite considerable theoretical work on serial order, it typically remains unexamined how physical agents may direct sequential actions at the environment within which they are embedded. Situated physical agents face a key problem - the need to accommodate variable amounts of time it takes to terminate each individual action within the sequence. Here we examine how Dynamic Field Theory (DFT), a neuronally grounded dynamical systems approach to embodied cognition, may address sequence learning and sequence generation. To demonstrate that the proposed DFT solution works with real and potentially noisy sensory systems as well as with real physical action systems, we implement the approach on a simple autonomous robot. We demonstrate how the robot acquires sequences from experiencing the associated sensory information and how the robot generates sequences based on visual information from its environment using low-level visual features.**

To accomplish even the simplest daily life tasks (e.g., brewing a cup of coffee or fixing a hole in a bicycle tire) requires people to plan and perform a series of actions in a particular order [1], [2]. To do so, people must keep their overall goal in mind and must remember the individual actions within the sequence which they need to direct at the pertinent perceptual objects. Critically, the actions must be initiated in the appropriate serial order. Since Lashley's classical work [3] there has been agreement that serial order is an important dimension of sequential behavior that may be dissociated from the capacity to remember the items and actions that occur in the sequence.

An aspect of the problem of sequential action that has not received much attention is that people are routinely able to generate behavioral sequences under conditions in which the duration of any individual action may vary greatly along the sequence, possibly in ways that are unpredictable beforehand. For instance, when making coffee, the time needed to fill the water container may be much longer than the time needed to close the lid of the coffee maker. In an assembly task, to mention another example, each gesture may take up different amounts of time depending on how precisely controlled the movement needs to be. Across repetitions of the same assembly task, some gestures may take up varying amounts of time depending on how successfully they are performed (think of slipping a small spring onto a hook when dealing with a clockwork). In daily life such variability is the rule rather than the exception. In contrast, many of the tasks through which serial order has been scientifically studied including musical performance, speech production, and typing are characterized by relatively uniform and predictable durations of the sequence elements. Accordingly, the stability of sequential action under such variable timing of the individual elements of a sequence has not been a major theme for most models of serial order. A related problem is how the progress along a sequence of actions is controlled by potentially noisy sensory information, itself of variable duration and quality.

Among models of sequence generation, those allowing for sensory information to intervene while the sequence is acted out provide the best chance to solve this problem [1]. Virtually all models of sequence generation, however, are based on relatively abstract representations both of the action systems that propel along the sequence as well as the sensory events that provide feedback about what happens in the world (for review see [2] and [4]). Stabilizing a sequence against variable timing in the presence of variable sensory feedback is critical to real action systems which impact on the environment from which they receive real-time sensory feedback. Based on such feedback, these systems must detect when one action has been terminated so as to autonomously transition to the next action.

This paper provides a new theoretical approach to sequence generation that, we believe, solves this stability problem in a principled fashion. The approach is based on Dynamic Field Theory (DFT), a neuronal grounded framework in which attractor dynamics and their instabilities generate motor, perceptual and cognitive function (review [5]). The neuronal representation of each state within the sequence is a stable state that is coupled into an action system. Sensory feedback about a condition-of-satisfaction [6] that signals that a step in the sequence has been successfully completed, is likewise represented by a stable neuronal state. This signal reliably brings about the switch from the previous to the next action.

We demonstrate that this approach can be acted out using real sensors and real effector systems by implementing the DFT model on an autonomous robot vehicle with a vision sensor. As a simple demonstration, the robot learns to sequentially search for an object of a given color. The robot learns a sequence of colors by being shown colored objects in the desired order. It is then capable of performing the sequential search task, in which the time needed to find an object of the color required at each step is unpredictable.

The DFT architecture builds on a spatial representation of serial order (a "positional" encoding in the classification of [7]). Our specific choice of representation is based on neurophysiological evidence for a neural encoding of serial order. For instance, when monkeys make sequences of pointing movements in response to a stimulus sequence, neurons in the anterior cingulate cortex fire specifically when the monkey is engaged in the action at a particular ordinal position (the second action, say), irrespective of which target that movement is directed to or which stimulus is used to cue that target [8], [9]. A similar encoding of ordinal position in a sequence is inferred from data that compare natural action sequences to stimulus induced goal-directed actions [10]. Neuronal pools responsive to serial information have also been found in motor cortex [11].

I. Dynamic Field Theory (DFT)

Originally an abstraction of the homogeneous neuroanatomy of many cortical and subcortical neural networks [12], DFT has become a framework within which neural process models can be generated for behaviors that reflect motor, perceptual and cognitive function [13]–[15]. Central to DFT is the assumption that these processes are characterized by continuous metric variables, encoded along the dimensions, $\mathbf{x}$, of neural activation fields, $u(\mathbf{x},t)$. These span the space of possible behavioral states, for instance, through movement parameters, the feature dimensions of perceptual representations, or the low-dimensional encoding of potentially high-dimensional memory representations [16]. The units of representation are localized peaks of activation which emerge as attractor solutions of the field dynamics, generically modelled as

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

where τ is a time constant, $h < 0$ the resting level, $S(x,t)$ an input function, $w(x-x')$ a "Mexican-hat" interaction function with short-range excitatory and long-range inhibitory connectivity, and $f(u)$ a sigmoidal nonlinearity [17]. Once sufficient positive activation has been induced near a field site, x, the local excitatory interaction stabilizes peak or "bump" solutions against decay while global inhibitory interaction stabilizes these solutions against diffusive spread. The self-stabilized peak solutions may emerge as stable states from an instability (the detection instability, see [5]), when localized input is increased in strength or when weak localized input is combined with a global, homogenous boost (modelled, for example, as an increase of h). Such peak solutions may persist in the absence of localized input (sustained activation) when the field dynamics is above the memory instability, at which activation levels become sufficiently high to engage neuronal interaction. Sustained peaks are a widely invoked model of working memory (review, [15], [18]). Long-term storage of metric information can be brought about by introducing inhomogeneities in the field through memory trace dynamics of various forms, from which localized peaks can be recreated through homogeneous boost [5].

II. The Dynamic Field Architecture for Sequence Generation

At the core of the Dynamic Field model of sequence generation lies a stack of neuronal activation fields (Fig. 1). All fields span the same dimensions which encode the feature values relevant to the performance of the actions. In the robotic implementation demonstrated below this feature dimension is the color of the visual targets that the robot must search in a learned serial order. Each layer in the stack represents a particular ordinal position in the sequence. Homogeneous connections among the fields in the stack control the sequential activation of localized peaks in one field at a time. Activation patterns in each field are preshaped by input from a layer of neurons that encode the learned sequences. Outside the stack, an additional output field defined over the same dimensions represents the currently activated action through a peak induced by the active ordinal field. The output field controls the action system, ultimately guiding the behavior of the autonomous robot. The output field also determines which pattern of sensory information in the robot's vision system signals the completion of the current action and provides input to a neuronal representation of a condition-of-satisfaction [6], which triggers the transition to the next ordinal position. An intuitive description of how this DFT architecture works follows below, while the mathematical equations are listed in the Appendix.

A. *Ordinal fields*

The stack of neural fields on the left in Fig. 1 encodes sequential order in that, aside from brief transitions, only one ordinal layer at a time can have a self-stabilized peak. The peak in an ordinal layer homogeneously inhibits the predecessor layer. Conversely, the peak homogeneously excites the successor layer. The excitatory coupling is shunted, however, by the condition-of-satisfaction system. When that system sends a signal to all ordinal fields in the stack, it activates the excitatory coupling only of the ordinal layer that sports a peak matching the current peak in the output layer (thus making sure, that the sensory signal reflects the outcome of a valid action within the sequence). The peak induced in the successor layer then suppresses the original peak in the predecessor layer. Where in the field the new peak is located is determined by prior activation induced by the long-term memory of the sequence.

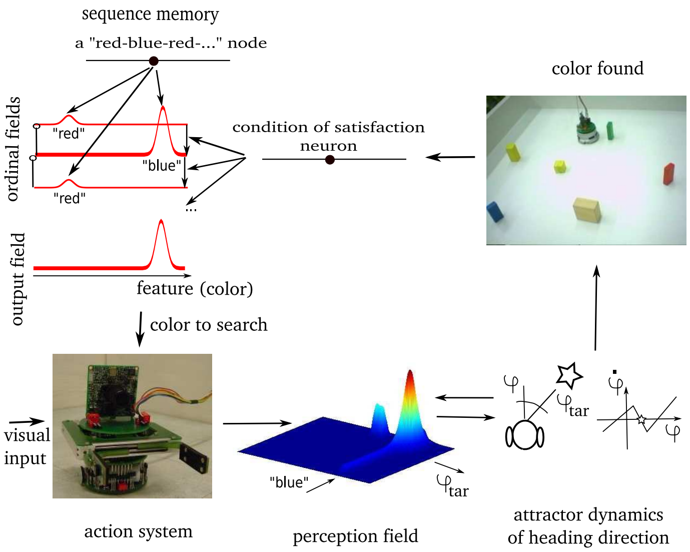

Fig. 1. Survey over the DFT architecture for generating action sequences.

B. Output field

The peak in the stack of ordinal fields is transmitted to the output layer, where it remains as a stable state until the switch to the next steps has been achieved. The output layer thus stably represents the current action throughout the variable time interval that the physical realization of this action takes. The output field also determines which sensory signal is sent to the condition-of-satisfaction system. This ensures, that the transition to the next state is caused by a sensory signal that is consistent with the current action.

C. Action system, perception field and heading direction dynamics

The action system may differ in different implementations of the DFT sequence generation architecture. In our robotic demonstration, the action system comprises the physical hardware, a perception field, an attractor dynamics of heading direction, and an associated robotic controller (bottom row of Fig. 1). We used a Khepera mobile robot equipped with an on-board color video camera. The perception field is a two-dimensional neuronal field that associates the color of visual targets with the heading direction, in which the targets are seen. It receives two-dimensional input extracted from the camera image: A histogram of hue values obtained within each column of the camera image defines the input function along the color dimension at the heading direction, into which this image column is pointing. A peak of activation in the output layer of the sequencing system provides a ridge of input across all heading directions, which effectively boosts all those parts of the visual array, where objects are seen with a color matching the current state of the output field. This leads to a self-stabilized peak in the two-dimensional field when there is such matching input. If there are several candidate objects in the visual array, the two-dimensional perception field selects one (typically the largest object) and then stabilizes that decision due to the dynamics (1).

The robot is controlled by a dynamics of heading direction which integrates two kinds of contributions [19]. The position of an activation peak in the perceptual field along the axis of heading direction controls an attractive force-let. Distance signals obtained from on-board active infra-red sensors modulate the strength of repellors, leading to obstacle avoidance. The rate of change of heading direction is input to servo-controllers on the robot vehicle, specifying the difference in velocity of the left and right active wheel so that the desired turning rate is generated. The forward velocity of the vehicle is also controlled so that it is slow when the robot is searching, is near an obstacle, or is close to a target. The velocity is faster when a target has been detected toward which the robot is moving. The forward velocity sets the mean of the signals sent to left and right wheel servo.

The combined effect of these dynamics is that the vehicle moves forward while avoiding obstacles, so that the system effectively searches its environment. As soon as the robot has detected an object with the currently requested color, it moves toward that object. The visual image of the object ultimately becomes sufficiently large within the camera plane to trigger a condition-of-satisfaction signal.

D. Condition-of-satisfaction system

A fundamental conflict of sequence generation is between the need to stabilize the behavioral state at a given step in the sequence and the need to destabilize that state in order to switch to the next step of the sequence. In the DFT architecture, this transition is mediated by a condition-of-satisfaction system (middle top panel of Fig. 1). This is a neuronal dynamics with the same bistability between “on” and “off” states as the neuronal field. In fact, the single neuron we are using could be viewed as the activation in a peak of a neuronal field, which could accommodate a range of conditions-of-satisfaction neurons needed in more complex scenarios. The condition-of-satisfaction neuron receives sensory input from the robot’s vision system. The count of pixels whose current color matches the color specified by the output field cues when the planned action has been achieved. When this input reaches a critical level, the “off” state of the condition-of-satisfaction neuron becomes unstable and the neuron switches to “on”. Because the sensory signal is automatically generated from real sensory input, its strength and duration fluctuates (bottom right of Fig. 4). To stabilize the condition-of-satisfaction system, the neuronal dynamics is bistable, so that the neuron remains “on” even if the sensory signal drops below the initial threshold. Only when the neuron is actively inhibited by negative input does it return to the “off” state. Such negative input comes from the stack of ordinal fields when it is in transition between two states: the input function detects correlation between suprathreshold activation in an ordinal field and in the output field, observed simultaneously for two consecutive ordinal layers.

The condition-of-satisfaction neuron shunts excitatory coupling from any ordinal field to its successor. Thus, the

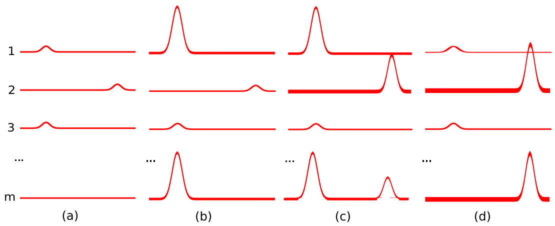

Fig. 2. Four stages of the sequence generation dynamics illustrated by showing the patterns of activation in three layers of the ordinal stack (1, 2, 3) and in the output layer (m). (a) Before sequence generation starts, the neural fields are preshaped through input from higher-level neuron, which represents the sequence as a whole. (b) A homogeneous boost to the first ordinal field starts sequence generation, inducing a localized peak of activation at the preshaped location. A matching peak is induced in the output field which drives and provides input to the action system. (c) Sensory feedback about the termination of the ongoing action shunts the homogeneous interaction among the sequence coding fields. This enables the first ordinal layer to induce a peak in the second ordinal layer at the preshaped location. (d) The newly established peak in the second ordinal layer inhibits the original peak in the first ordinal layer and resets the output field so that it carries a peak at a matching location.

condition-of-satisfaction system does not need to "know" at which ordinal position in the sequence it currently is.

E. Sequence memory and preshape of the ordinal fields

Laying down a memory trace of self-stabilized peaks of activation is a simple form of learning in DFT [13]. The memory trace preshapes the activation field, so that, conversely, a localized peak can be induced from the low level of preactivation by a homogeneous boost. (A possible neuronal realization of this learning mechanism is a Hebbian strengthening of local excitatory connections, which will likewise support peak induction from homogeneous input.). Based on this mechanism, sequence learning is realized by accumulating a memory trace during a training session (see below). A long-term memory of the memory trace is stored in the synaptic connections from a higher-level pool of neurons to the ordinal stack (top left of Fig. 1). Different members of the pool encode different sequences. To enact a particular sequence, the corresponding neuron in the sequence memory pool is activated, so that field in the ordinal stack are preshaped at the memorized locations.

Figure 2 illustrates the stable states and their instabilities that are realized when a sequence is initiated, the first action is realized and the transition to the second action occurs.

III. RESULTS

A. Learning

First, we demonstrate one-shot learning of a sequence (Fig. 3). Five colored bricks are presented to the robot in succession. In each case, input from the camera generates a peak in the perceptual field at the corresponding hue value and heading direction. The supra-threshold component of the peak summed along heading directions is a distribution of activation over hue space and is fed into each field in the ordinal stack. When the first brick is presented, the first ordinal field receives a homogenous boost. This causes the first layer to develop a peak at the color specified by the input from the perceptual field. A memory trace associated with each ordinal layer accumulates activation at the peak location. This memory trace subsequently controls the autonomous acquisition of the sequence. When the memory trace reaches a critical mass and the perceptual field registers a removal of input (because the brick has been removed from the visual array of the robot), a signal is generated to the condition-of-satisfaction neuron which now autonomously organizes the switch to the second ordinal layer. When the second brick is presented, this layer develops a peak at the associated hue value, suppressing the peak in the first layer and laying down a memory trace. This continues until all colored objects have been shown to the agent. The pattern of memory traces is associated with a neuron in the long-term sequence memory (dynamics of this process is not modelled here).

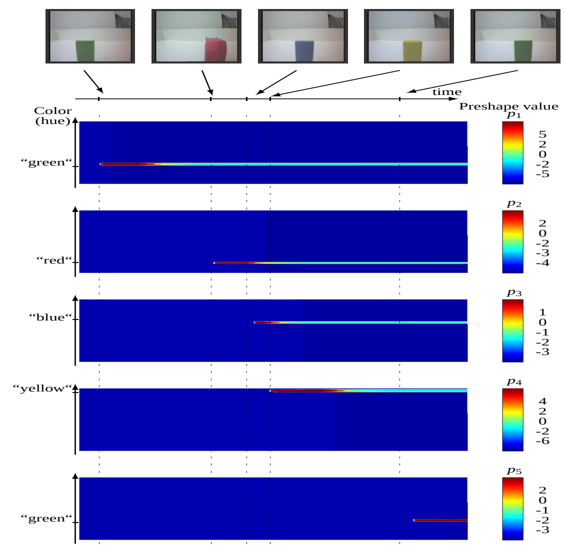

Fig. 3. The top row illustrates a time line of visual stimulation from left to right. Below are shown the time-courses for dynamic preshapes of five neural fields in the ordinal stack, which end up sequentially encoding the five colors in the order in which they were presented.

B. Sequence generation

To execute the sequential search task, the robot is set into an arena in which colored bricks have been distributed (top right of Fig. 1). The "go" signal is given, leading to activation of a peak in the first layer of the ordinal stack at the location representing the first color in the sequence and to activation of an associated peak in the output field (Fig. 4). This peak provides a ridge input at the corresponding hue value to the two-dimensional color-heading direction perceptual field (bottom of Fig. 1). The vehicle moves around the arena, avoiding obstacles. The vehicle keeps track of the heading direction by integrating its rate of change. The camera image is continuously transformed into the coordinate frame of the heading direction and input into the perceptual field. When visual input is encountered that overlaps sufficiently with the preshaped ridge at the requested hue value, a peak forms in the perceptual field. This peak defines a movement target and the vehicle moves in the required direction, while continuing to avoid obstacles. Typically, the vehicle is able to approach the object, so that the object's projection onto the visual array grows in size, until it looms sufficiently large in the camera image to send a sensory signal to the condition-of-satisfaction system (bottom right of Fig. 4).

At this point, a transition in the ordinal stack is triggered. The second ordinal layer builds a peak at the second learned color value, which is replicated in the output layer. As this happens, the condition-of-satisfaction neuron switches back to the "off" state and the peak in the first ordinal layer is suppressed by inhibition from the second ordinal layer. The change of location of the peak in the output layer shifts the hue value at which a ridge is input into the perceptual field. The current peak in this field therefore decays. This removes the attractor of heading direction in the direction of the previous target and effectively puts the system back into search mode. The robot again moves around, avoiding obstacles, until it encounters in its visual array enough color information that overlaps with the new preshaped ridge, leading to the generation of a new peak in the perceptual system. This reinstates a movement target, toward which the robot navigates. As that target is approached, another sensory signal is sent to the condition-of-satisfaction system, leading to the second transition.

Figure 4 shows four such transitions as recorded from the life robot. In each case, the peak in the output layer lasts as long as it takes to reach a target of the specified color. Naturally, the time needed to find a target varies depending on the configuration of the robot and the environment. The core feature of the DFT approach to sequence generation is illustrated here: The system operates stably in the face of such variable and unpredictable timing of each individual action. Moreover, the sensory signal sent to the condition-of-satisfaction system is obtained from the life camera image on the robot. The movement of the robot as well as intrinsic properties of the video system make this a noisy signal, but the bistable condition-of-satisfaction system stabilizes the decision that the action goal has been reached, leading to orderly transitions. Note also, how the duration of the transition itself varies, as reflected in the time interval during which the condition-of-satisfaction neuron is "on". This reflects the speed of the transition in the ordinal stack, which depends on the metric distance between the successive peaks, the strengths of the learned patterns, and on fluctuations in the neuronal dynamics.

IV. Discussion

We have introduced a neurally grounded architecture for the learning and generation of behavioral sequences based on the framework of Dynamic Field Theory. To demonstrate the system's capability to tolerate variable durations of each action within a sequence, we implemented the architecture on an autonomous robot that searches for colored objects in a learned order of colors. The ordinal position of an action is encoded along a stack of neuronal activation fields, each of which expands feature dimensions needed to specify actions. This feature dimension thus represents the "contents" at each step, a form of positional encoding in the classification of [20]. While we have implemented a single feature dimension, color, in our robotic example, these fields may link to rich representations that bind multiple different feature maps into perceptual objects [16]. By steering the perceptual system, this representation makes it possible to generate actions directed at objects in the world as illustrated in our robotic demonstration.

Like ours, the model of [21] is based on evidence for neurons encoding ordinal position. These authors show how short-term memory for sequences can be implemented in an auto-associative neural network with attractor dynamics. Although the model is couched in terms of putative neuronal mechanisms, its functional dynamics can be described as a sequence of semi-stable states. The transitions are triggered internally through a global inhibitory input and a synaptic adaptation mechanisms, not by sensory input from the world. The system does not, therefore, tolerate variable durations of each action step. The system has not been implemented on an autonomous robot. It is not clear, if it can generate actions oriented at perceptual objects.

Botvinick and Plaut present a recurrent neuronal network which has a similar mission as ours, generating behavioral sequences [1]. Their network model is essentially a neuronal dynamics, although time is treated in discrete steps. Like ours, this model is in principle linked to sensory input from the world, so that it may direct actions at objects in the world. The model is far from a real-world implementation, however, and does not address the question of how actions of variable durations may be accommodated. A number of similar neuronal network models in the positional encoding framework are directed at generating sequences at the level of abstract representations and do not pose the question of how such systems may deal with the real-time control of effectors or with online perceptual information [2], [4], [22]. Unlike these more abstract models, we have not yet addressed serial

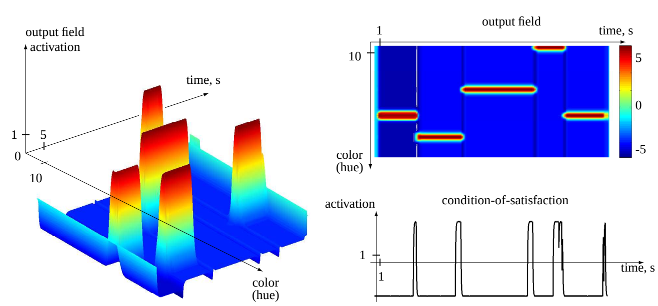

Fig. 4. The temporal dynamics of the output field during execution of the sequence "green-red-blue-yellow-green". Left: Self-stabilized peaks in the output field represent the sequential color values that control the search behavior of the robot. Right: A color-coded two-dimensional view of the temporal evolution of the output field is aligned with the time series of the activation level of the condition-of-satisfaction neuron. Events at which the signal is positive indicate the points in time when the system transitions to the next stable state within the sequence.

order errors.

The architecture presented here has a few other obvious limitations. Extension is easy in some cases, such as increasing the number of sequences learned, adressing hierarchical action sequences, autonomously organizing sequential memories, and improving memory representations across multiple trials. Making sequence generation more flexible, so that perceptual experience at one step may affect the continuation of the ongoing sequence, may provide a key to a more autonomous form of exploration and learning of behavioral sequences.

V. Appendix: Mathematical description of the model

The dynamics of each field, $u_i(x,t)$ in the ordinal stack (ordinal index, $i = 1,\ldots,N_{Sc}$) is

$$\begin{aligned}\tau_{\mathrm{Sc}}\dot{u}_i(t,x) &= -u_i(t,x) + h + \int f(u_i(t,x'))w_{\mathrm{Sc,Sc}}(x,x')dx' \\ &+ C_+ f(\xi_{\mathrm{cs}})\int f(v(x',t))f(u_{i-1}(x',t))dx' \\ &- C_- \int f(u_{i+1}(x',t))dx' + P_{iY}(x,t)\end{aligned} \tag{2}$$

with similar parameters as the generic field equation Eq. 1. The constants, C_+ and C_- control the boosting and deboosting coupling, the preshaping input, $P_{iY}(x,t)$ comes from the higher level neuron Y. Sensory feedback about action completion comes through the condition-of-satisfaction neuron, ξ_{cs}.

The dynamics of the motor field $v(x,t)$ is:

$$\begin{aligned}\tau_M\dot{v}(x,t) &= -v(x,t) + h + \int f(v(x',t))w_{MM}(x,x')dx' \\ &+ \Sigma_{i=0}^{N_{Sc}}\left(\int f(u_i(x',t))w_{MSc}(x,x')dx' + C_+\right)\end{aligned} \tag{3}$$

with analogous notation.

The perceptual field, $u_{\mathrm{p}}(x,\phi)$ is a 2D variant of the Amari system, Eq. 1. The dynamics of heading direction, ϕ is

$$\tau_\phi\dot{\phi} = \lambda_{\mathrm{obs}}F_{\mathrm{obs}}(\phi) - \lambda_{\mathrm{tar}}\int f(u_{\mathrm{p}}(x,\psi))(\psi-\phi)dxd\psi \tag{4}$$

The obstacle forcelet $F_{\mathrm{obs}}(\phi)$ is a sum over the contributions of the 6 infrared sensors of the Khepera robot. Each sensor contributes with strength depending on the sensed distance to an obstacle, and angular range depending on the opening angle of the sensors, on the robot size, and on the sensed distance (see [23] for details).

The condition of satisfaction neural dynamics is

$$\dot{\xi}_{\mathrm{cs}} = -\lambda\xi_{\mathrm{cs}} + h_{\mathrm{cs}} + \mu\ f(\xi_{\mathrm{cs}}) + I(t) + F(t) \tag{5}$$

where λ and μ are constants, h_{cs} is the resting level, and $f(\cdot)$ is the sigmoidal non-linearity, which provides self-excitation to the neuron, ξ_{cs}. The sensory signal, $I(t)$, is obtained from the vision system, the negative reset signal, $F(t)$, is obtained from the ordinal stack, both described in the main text.

The memory trace mechanism for learning a sequence is implemented as:

$$\begin{aligned}\tau_p\dot{p}_i(x,y,t) &= (\lambda_{build}f(u_i(x,t))\cdot(-p_i(x,y,t) \\ &+ f(u_i(x,t)) * f(U(y,t))),\end{aligned} \tag{6}$$

where $p_i(x,y,t)$ is a synaptic connection from the site x of a sequence coding field $u_i(x,t)$ to a site y in a higher level pool of neurons, $U(y,t)$, τ_p is the time constant, λ_{build} is the rate of strengthening of connections driven by simultaneous activation of $u_i(x,t)$ and $U(y,t)$.

References

[1] M. Botvinick and D. C. Plaut, "Doing without schema hierarchies: A recurrent connectionist approach to normal and impaired routine sequential action," *Psychological Review*, vol. 111, no. 2, pp. 395–429, 2004.

[2] G. Houghton and T. Hartley, "Parallel models of serial behaviour: Lashley revisited," *Psyche*, vol. 2, no. 25, 1995.

[3] K. Lashley, "The problem of serial order in behavior," in *Cerebral mechanisms in behavior: The Hixon symposium(pp. 112 - 146)*, 1951.

[4] R. N. Henson, "Short-term memory for serial order: The start-end model," *Cognitive Psychology*, no. 36, pp. 73–137, 1998.

[5] G. Schöner, "Dynamical systems approaches to cognition," in *Cambridge Handbook of Computational Cognitive Modeling*, R. Sun, Ed. Cambridge, UK: Cambridge University Press, 2008.

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

[7] R. N. A. Henson and N. Burgess, "Representations of serial order," in *Connectionist Representations*, J. A. Bullinaria, D. W. Glasspool, and G. Houghton, Eds. Springer Verlag, 1997, pp. 283–300.

[8] E. Procyk, Y. Tanaka, and J. Joseph, "Anterior cingulate activity during routine and non-routine sequential behaviors in macaques," *Nature Neuroscience*, vol. 3(5), pp. 502–508, 2000.

[9] E. Procyk and J. Joseph, "Characterization of serial order encoding in the monkey anterior cingulate sulcul," *European Journal of Neuroscience*, vol. 14, pp. 1041–1046, 2001.

[10] J. W. Aldridge and K. C. Berridge, "Coding of serial order by neostriatal neurons : A "natural action" approach to movement sequence," *The Journal of Neuroscience*, vol. 18(7), pp. 2777–2787, 1998.

[11] A. F. Carpenter, A. P. Georogopoulos, and G. Pellizzer, "Motor cortical encoding of serial order in a context-recall task," *Science (Reports)*, vol. 283, pp. 1752–1757, 1999.

[12] H. Wilson and J. Cowan, "A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue," *Biological Cybernetics*, vol. 13, pp. 55–80, 1973.

[13] W. Erlhagen and G. Schöner, "Dynamic field theory of movement preparation," *Psychological Review*, vol. 109, pp. 545–572, 2002.

[14] E. Thelen, G. Schöner, C. Scheier, and L. Smith, "The dynamics of embodiment: A field theory of infant perseverative reaching." *Brain and Behavioral Sciences*, vol. 24, pp. 1–33, 2001.

[15] J. P. Spencer and G. Schöner, "Bridging the representational gap in the dynamical systems approach to development," *Developmental Science*, vol. 6, pp. 392–412, 2003.

[16] J. S. Johnson, J. P. Spencer, and G. Schöner, "Moving to higher ground: The dynamic field theory and the dynamics of visual cognition," *New Ideas in Psychology*, 2007.

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

[18] D. J. Amit, "The hebbian paradigm reintegrated: Local reverberations as internal representations," *Behavioral and Brain Sciences*, vol. 18, no. 4, pp. 617–626, 1994. [Online]. Available: http://www.bbsonline.org/documents/a/00/00/04/21/bbs00000421-00/bbs.amit.html

[19] G. Schöner, M. Dose, and C. Engels, "Dynamics of behavior: Theory and applications for autonomous robot architectures," *Robotics and Autonomous Systems*, vol. 16, pp. 213–245, 1995.

[20] R. Henson and N. Burgess, "Representations of serial order," 1997.

[21] G. Deco and E. T. Rolls, "Sequential memory: A putative neural and synaptic dynamical mechanism," *Journal of Cognitive Neuroscience*, vol. 17:2, pp. 294–307, 2005.

[22] M. M. Botvinick and D. C. Plaut, "Short-term memory for serial order: a recurrent neural network model." *Psychological Review*, vol. 113, no. 2, pp. 201–233, 2006.

[23] E. Bicho, P. Mallet, and G. Schöner, "Target representation on an autonomous vehicle with low-level sensors," *The International Journal of Robotics Research*, vol. 19, pp. 424–447, 2000.