Adaptive Behavior, 19(5), pp.295-316, 2011

Simultaneously Emerging Braitenberg Codes and Compositionality

Shorttile: Emerging Braitenberg Codes and Compositionality

Yuuya Sugita, Jun Tani, Martin V. Butz*
*corresponding author
Department of Psychology III
University of Würzburg
Röntgenring 11, 97070 Würzburg, Germany,
Tel.: +49 931 31 82808
Fax: +49 931 31 82815
Email: butz@psychologie.uni-wuerzburg.de

Simultaneously Emerging Braitenberg Codes and Compositionality

Yuuya Sugita
Department of Psychology III (Cognitive Psychology)
University of Würzburg
Röntgenring 11, 97070 Würzburg, Germany
Email: yuuya.sugita@gmail.com

Jun Tani
RIKEN Brain Science Institute
2-1 Hirosawa, Wako-shi, Saitama 3510198, Japan
Email: tani@brain.riken.jp

Martin V. Butz
Department of Psychology III (Cognitive Psychology)
University of Würzburg
Röntgenring 11, 97070 Würzburg, Germany
Email: butz@psychologie.uni-wuerzburg.de

Abstract

Although many researchers have suggested that compositional concepts should be sensorimotor grounded, the method to accomplish this remains unclear. This paper introduces a second-order neural network with parametric biases (sNNPB) that learns compositional structures based on sensorimotor time series data. The data was produced by a simulated robot that executed distinct object interactions (move-to and orient-toward). We show that various sNNPB setups can learn to compositionally imitate object-interactions beyond the interactions that were specifically trained, which was not possible with previous neural network (NN) architectures, including recurrent neural networks (RNNs). We also show that these imitation capabilities are accomplished by developing a self-organized, geometrically-arranged compositional concept structure in the PB values and task-oriented, Braitenberg-like sensory encodings in hidden sensory layers. Since second-order connections were necessary to accomplish the task, we hypothesize that such connections may be essential to drive the learning of both sensorimotor-grounded compositional structures and Braitenberg-like, behavior-oriented "pro-presentations". From a cognitive perspective, we show how sensorimotor time series of interactions may be processed to generate the signals necessary to develop semantically compositional structures.

Keywords: the origin of compositionality, Braitenberg codes, sensorimotor similarity, constructivist approach, connectionist network, cognitive modelling

Introduction

Compositional concept structures enable the meaningful manipulation, combination, and relation of concepts. For the involved processes to become meaningful, structured *concept elements* need to be associated with (a) role-governed (Markman & Stilwell, 2001) or slot-filler *categories* (Nelson, 1988) and (b) *role-argument* structures (Barsalou, 1999), which determine the possible proper combinations of elemental concepts. To foster the development of such meaningful concepts, Karmiloff-Smith (1992) proposed that the necessary element-category-argument structures need to be derived from implicit, sensorimotor representations – unlike nativist models, which assume innate compositional representations (Fodor & Lepore, 2002, cf. e.g.). With her representational rewriting (RR) model, Karmiloff-Smith (1992) suggested that compositional concepts must be acquired through the "re-writing" of pre-acquired holistic sensorimotor concepts. The underlying computational processes, however, have not been clarified yet.

In the investigation of pre-linguistic concept structures, namely image schemas, Johnson (1987) and Lakoff (1987) pointed out that semantic categories and analogies appear to be closely related to principles of sensorimotor interactions and behavioral control. Most recently, Pastra and Aloimonos (2011) even pointed out that the human cognitive apparatus contains a minimalist grammar for action, which has many properties of the universal grammar of Noam Chomsky, and which appears to be derived from our behavioral control system. However, also in this case, the computational processes that underlie the sensorimotor-based formation of image schemas have not been determined, yet. It is particularly puzzling how image schemas may be derived from sensorimotor experience but meanwhile may be a prerequisite for organizing these experiences (Clausner, 2005). An attempt to implement an RR model based on image schemas (Mandler, 1992) suffers from essentially this difficulty. Consequently, a principle to organize sensorimotor experiences may be required if a compositional system is to be acquired based on sensorimotor experiences (Mandler, 1992, cf.).

We investigate how a neural network architecture may learn to relate holistic concepts to each other, thus developing compositional concept structures. We hypothesize that the embodiment of the agent with its continuous sensorimotor dynamics, including the involved learning mechanisms, can derive a compositional system of concepts based on *sensory and motor similarities*. The idea is based on the *structural alignment hypothesis* (Gentner & Markman, 1997), which proposes a continuous developmental path from similarity-based to analogy-based transfer in problem solving. While analogy-based transfer usually refers to knowledge transfer across domains, such as between fluid dynamics and electromagnetics, it is also applicable to compositionality. Gentner and Markman (1997) invented a common computational process underlying both types of transfer, but their approach relies on pre-defined compositional representations. In our work, structural alignment is reconsidered from a constructivist point of view, replacing analytical similarity with skill transfer-based similarity, which is identified by our learning architecture in sensorimotor time series data.

To achieve this task, we introduce a second-order neural network with parametric biases (sNNPB). The sNNPB is trained to imitate particular object interactions, which are generated with a simulated robot platform. During training, no semantic, compositional indicators are provided. We show that the sNNPB develops compositional, geometrically arranged structures based on the sensorimotor time series data produced by the robot platform. The network develops these compositional codes solely due to the training setup, the modular sNNPB architecture, and the sensory- and motor-capabilities of the simulated robot.

The sNNPB is partitioned into a visually-driven sensory-to-motor sub-network (*sem-net*) and a PB-driven sub-network (*pb-net*). The *pb-net* determines the mapping (thus second-order network) from the *sem-net* to the motor output (cf. Fig.3). Due to this setup, the *sem-net* develops behavior-oriented, distributed Braitenberg-like encodings (Braitenberg, 1984). The *pb-net* develops compositional structures that activate those sensory-to-motor mappings that result in the desired object interaction. A detailed structural analysis of the resulting network structure reveals that the sNNPB learned to embed semantically compositional structures in self-organized, geometric manifolds.

In contrast to previous multitask learning approaches (Thrun & O'Sullivan, 1996; Caruana, 1997), our sNNPB architecture learns to categorize time series data by concurrently self-organizing both (a) a sensory input space for guiding dynamic sensorimotor flow and (b) a parametric bias space for the compositional activation of particular interactions. Seeing that previous NN architectures including RNNPBs (Tani, Ito, & Sugita, 2004) and other RNNs could not develop similar behavioral generalization patterns, we conclude that multiplicative, second-order network architectures may be essential to develop semantically compositional structures. In sum, this study devises a potential continuous developmental path from similarity to compositionality. Meanwhile, the study provides indicators that the separation of motor control in a dorsal, control-oriented and a ventral, decision-oriented path may be a prerequisite to develop compositionality and, ultimately, language.

We now first introduce the simulated robot architecture, the neural network architecture, and the different settings evaluated. We then provide a detailed performance and structural network analysis, revealing how compositionality is represented and how the task is solved. Finally, we conclude the paper with a general discussion.

Robot System Setup

Our robot setup is based on a psychological experiment conducted by Meltzoff (1988b), who examined the conceptualization ability of pre-linguistic infants. The original study reported that 9-month-old infants could imitate previously unseen object-manipulating actions for 24 hours after the presentation of the actions. This indicates that pre-linguistic infants memorize presented actions in an abstracted representation, which may provide a foundation for later language acquisition (Meltzoff, 1988a).

In our experiment, the setup is modified in order to focus on the emergence of compositionality. Interactions are presented by steering a simulated robot by means of a teaching program. The robot perceives its own visual input and motor output during such interaction episodes, omitting the necessary translation between allocentric and egocentric views during imitation. However, the time scale of the experiment is expanded because the robot needs to develop structured representations for encoding the encountered interactions compactly in role-argument frames.

Robot and Environment

The simulated robot consisted of a simple two-wheeled robot platform equipped with visual surround sensors. The body of the robot was a cylinder of 24cm diameter and 14cm height. The visual sensors were located at the center of the body (in simulation) and partitioned the covered 120° view into nine uniform areas for which the sensors reported the dominant color and the size of the sub-area that is colored.

In each interaction episode, the robot interacted with one of six colored cylinders (`blue`, `cyan`, `green`, `yellow`, `orange`, or `magenta`) that were 13cm in diameter and 25cm in height. Ei-

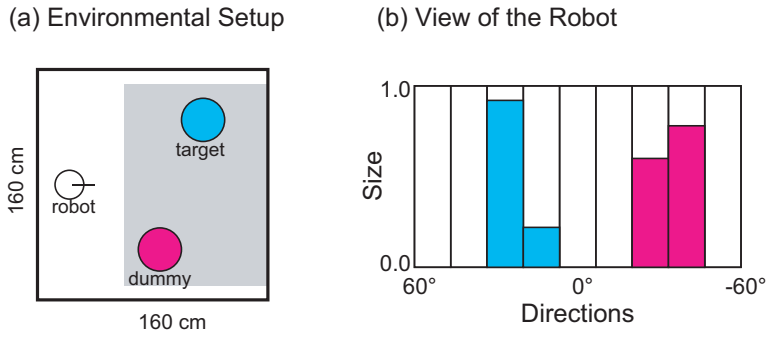

Figure 1. The robot starts with a fixed starting position. A target and an optional colored dummy object are placed randomly within the shaded area (W120cm × D134cm) shown in (a). The robot takes a visual input as shown in (b) in this environmental situation.

ther one or two randomly placed objects were present. Fig. 1 shows a possible situation and the corresponding sensory information. The interactions $\mathcal{I}$ are classified into 36 categories: two interactions: `moveto` (move to) or `orientto` (orient toward) one of the six targets, where the `orientto` interaction was further modified with five possible angular offsets: (`-30`, `-18`, `0`, `+18`, and `+30`). In the `moveto` interactions, the robot moves toward the target and stops just before touching it. In the `orientto` interaction, the robot orients itself toward the target with a certain offset by rotating its body. In the following sections, an interaction is denoted by a concatenation of labels of one of the two types, one of the six targets, and, optionally, one of the five offset angles; for example, `orientto-blue+18`. It should be noted that the labels are used for convenience only – the robot is never provided with any information about the semantic structure or compositionality of the interactions.

Four independent blocks of learning experiments were conducted with supervised data of different sparseness, as shown in Fig. 2. The subset of interactions on which the system is trained is denoted by $\mathcal{I}'$. We trained the sNNPB on these six particular subsets of interactions in independent experiments to illustrate the conceptualization capabilities, including similarity and compositionality biases. Each experimental epoch consisted of three phases: creation of training data, learning, and evaluating the imitation performance on all possible interaction types in a random subset of the possible scenarios.

Generation of Training Data

To gather sensorimotor time series data, the robot was controlled by a training program. For each goal-oriented behavioral interaction, 120 training sequences were generated in different environmental settings. In 20 out of the 120 cases only a target object was present, while in the remaining 100 cases a dummy object and a target object were located in the environment. The dummy object was chosen from the five objects (colors) other than the target, and 20 examples were recorded for each object.

The training program generated wheel speeds (v_L[cm/step], v_R[cm/step]) appropriate for executing a particular interaction, taking as input the distance r[cm] to and the direction θ[rad] to the target object. Formally, the program taught `moveto` interactions by moving the robot center to a

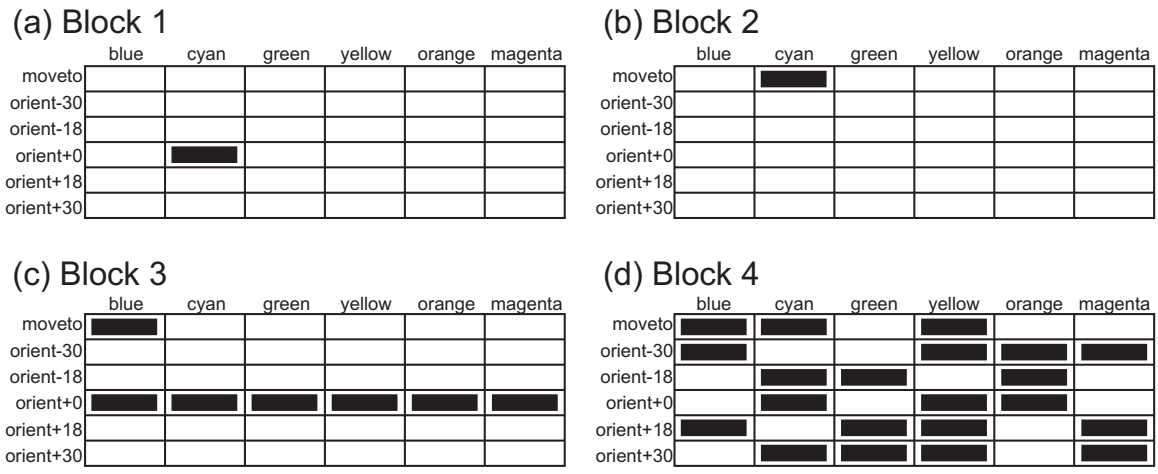

Figure 2. Four particular distributions of trained sets of interactions (blocks) are used to illustrate the sNNPB capabilities.

distance of about 35.0cm from the target object center as follows:

$$v_R = \max(-1.0, \min(v_\theta + v_r, 1.0)) \tag{1}$$

$$v_L = \max(-1.0, \min(-v_\theta + v_r, 1.0)) \tag{2}$$

$$v_\theta = \begin{cases} 0.25\theta & \text{if } r \geq 36.0, \\ 0.75\theta & \text{otherwise,} \end{cases} \tag{3}$$

$$v_r = \begin{cases} 0.0045r + 0.158 & \text{if } r \geq 36.0, \\ 0.0045(r - 35.0) & \text{otherwise.} \end{cases} \tag{4}$$

The program moves the robot forward with decreasing speed (by modifying v_r), at the same time orienting the robot towards the target object by modifying v_θ. It stops once it reaches a distance of less than 36.0cm to the target object. A typical `moveto` interaction episode consequently consisted of about 70 time steps on average.

Similarly, the program taught `orientto` interactions with offset angle ϕ as follows:

$$v_R = \max(-1.0, \min(0.5(\theta - \phi), 1.0)), \tag{5}$$

$$v_L = -v_R, \tag{6}$$

turning the wheels with opposite speeds to rotate the robot without forward movement. A typical `oriento` interaction episode consisted of about 20 time steps on average.

While the teaching program used precise state information, the sNNPB architecture received only the simulated visual information specified above and the generated wheel speeds. Note that both, the `orientto` and the `moveto` interactions translate the target angle into according positive and negative wheel speeds to realize a rotation towards the target object. However, the `orientto` command uses a different factor for this transition (0.5 versus 0.25) and the `moveto` command is folded into the forward movement. We will see, nonetheless, that the sNNPB is able to detect this time series compositionality, yielding (asymmetric) generalization capabilities.

Second-Order NN with Parametric Biases

We report experiments with three sNNPB setups below. We first specify the network architecture common to all three setups and then specify the distinctions between the different setups.

Network Architecture

The architecture of the implemented sNNPB is shown in Fig. 3. The network consists of two interacting sub-networks: a feed-forward sensor-to-motor network (*sem-net*) and a pb-based network (*pb-net*).

The *sem-net* is depicted on the left side of the figure. It takes as input the visual information at the vision input (VI) layer. Next, it transforms this information via a hidden layer (H) to the processed visual representation layer (VR) in the conventional feed-forward NN way:

$$\mathrm{VR}_i(t) = f_{\mathrm{VR}}\left(\sum_{j=0}^{N_{\mathrm{H}}} w_{\mathrm{VR}(i)\leftarrow\mathrm{H}(j)} \cdot \mathrm{H}_j(t)\right), \quad \mathrm{H}_i(t) = f_{\mathrm{H}}\left(\sum_{j=0}^{N_{\mathrm{VI}}} w_{\mathrm{H}(i)\leftarrow\mathrm{VI}(j)} \cdot \mathrm{VI}_j(t)\right), \tag{7}$$

where X_i ($X \in \{\mathrm{MO}, \mathrm{VR}, \mathrm{H}, \mathrm{VI}\}$) denotes the activity of a node, $w_{X_i \leftarrow Y_j}$ the connection weights, t the time stamp, and f_X the transformation function (a hyperbolic tangent or a linear function).

Finally, the information contained in VR is transferred to the motor output layer (MO) by:

$$\mathrm{MO}_i(t) = \sum_{j=0}^{N_{\mathrm{VR}}} \mathrm{sC}_{i,j} \cdot \mathrm{VR}_j(t), \tag{8}$$

where $\mathrm{sC}_{i,j}$ refers to the weights generated by the *pb-net* as specified below. The *sem-net* is thus a conventional layered feed-forward neural network, except that it has second-order connections (Pollack, 1990) between the VR and MO layers, meaning that the weights between VR and MO are flexibly set to particular values. This enables the *sem-net* to generate different sensorimotor interactions, dependent on the currently activated weights from VR to MO.

The weights of the second-order connections are determined by the *pb-net*, shown in gray on the right side of Fig. 3. The *pb-net* itself is also a conventional feed-forward neural network. It takes as input a concept vector of parametric biases (PB) and generates the weights of the second-order connections by means of the activity in the second-order connectivity (sC) layer, depicted as consisting of two sub-layers, which determine the weights for the two output nodes, respectively.

$$\mathrm{sC}_{i,j} = f_{\mathrm{sC}}\left(\sum_{k=0}^{N_{\mathrm{PB}}} w_{\mathrm{sC}(i,j)\leftarrow\mathrm{PB}(k)} \cdot \mathrm{PB}_k\right), \tag{9}$$

where sC_{ij} is the output of a node in the sC layer, which corresponds to a connectivity from the j-th VR node to the i-th MO node.

The generation of action outputs works as follows. First, a *concept vector* $\mathbf{u}_i$ needs to be available, which is used to set the PB node activities. This vector determines the values of the sC connection weights according to (9). These weights are set once at the beginning of an interaction episode and are kept fixed while a particular interaction unfolds. To generate actual motor outputs, sensory activities are transferred into motor output activities according to (7 and 8). Note that the time scale of the *pb-net* is thus different from the *sem-net*. The 'normal' connection weights $W = \{w_{\mathrm{H}\leftarrow\mathrm{VI}}, w_{\mathrm{VR}\leftarrow\mathrm{H}}, w_{\mathrm{sC}\leftarrow\mathrm{PB}}\}$ capture the common characteristics among all of the trained interactions $\mathcal{I}'$, whereas each concept vector $\mathbf{u}_i$ determines particular second order connections sC to realize particular object interactions $i \in \mathcal{I}'$.

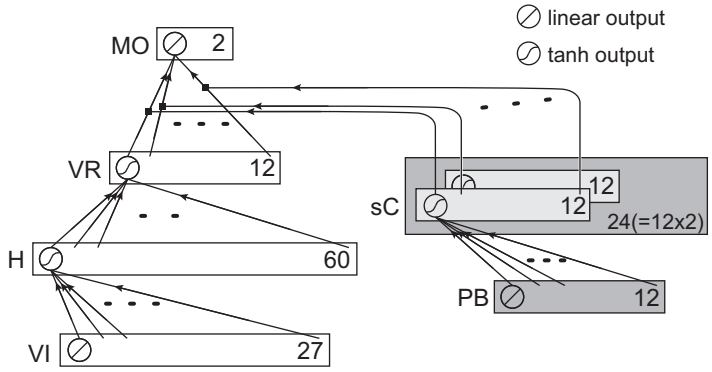

Figure 3. **The sNNPB architecture:** Each rectangle represents a layer, which contains multiple nodes. The number of nodes and the utilized output function are indicated in each rectangle. The *pb-net*, which consists of the gray layers, controls the connectivity from the VR to the MO layer. The second order connection layer sC specifies two subsets of weights, which connect the VR layer with the two motor output nodes of the MO layer.

Learning

Unlike a standard layered neural network, both connection weights W and interaction-respective concept vectors $\mathbf{u}_i$ ($i \in \mathcal{I}'$) are optimized in the sNNPB, where the vectors are stored in $U = \{\mathbf{u}_i \mid i \in \mathcal{I}'\}$. Both sets of parameters are optimized simultaneously by means of the conventional steepest descent method with respect to the output error defined as follows:

$$E(W^T, U^T) = \sum_{i \in \mathcal{I}'} E_i(W^T, \mathbf{u}_i^T) \tag{10}$$

$$E_i(W^T, \mathbf{u}_i^T) = \sum_{j=0}^{N_i} \sum_{t=0}^{l_{ij}} E_{ij}(t; W^T, \mathbf{u}_i^T) \tag{11}$$

$$E_{ij}(t; W^T, \mathbf{u}_i^T) = \|\hat{MO}_{i,j}(t) - MO(t; W^T, \mathbf{u}_i^T)\|^2, \tag{12}$$

where T denotes the current learning iteration, W^T the current connection weights, U^T the set of interaction concept vectors, $\mathbf{u}_i^T (\in U)$ a particular interaction concept vector to be optimized, $N_i (= 120)$ the number of pre-recorded training examples of the i-th interaction concept, l_{ij} the length of the j-th time series example of the i-th interaction, $\hat{MO}_{i,j}(t)$ the desired motor activity of the time series with respect to its corresponding visual input $\hat{VI}_{i,j}(t)$, and $MO(t; W^T, \mathbf{u}_i^T)$ the actual output of the network at that time.

The learning procedure is implemented by using the conventional back-propagation algorithm. All connection weights W^0 are initialized with uniformly distributed random values $\in [-0.1, 0.1]$ and all entries in each $\mathbf{u}_i^0 \in U^0$ are initially set to zero. Each learning example indicates its particular interaction concept correspondence i, which leads to the re-application of the corresponding current concept vector $\mathbf{u}_i^T$ and the according adjustment of that vector. Learning errors are back-propagated to their corresponding concept vectors $\mathbf{u}_i$, which are maintained and progressively adjusted throughout learning. In each of the reported experiments below, we conducted $30,000$ learning iterations. Alg. 1 specifies the precise learning procedure.

Algorithm 1 Learning procedure in sNNPB.

for all interactions i in $\mathcal{I}'$ **do**

Load the stored $\mathbf{u}_i^T$ to the PB layer.

for all pre-recorded examples j of the interaction i **do**

Calculate the delta errors of connection weights $-\partial E_{ij}/\partial W$ and of PB vector $-\partial E_{ij}/\partial \mathbf{u}_i$ by using the back-propagation algorithm

end for

Sum up the delta errors over all time steps t of all time series (i, j) to obtain $\delta\mathbf{u}_i^{T+1}$.

$$\left(\because \delta\mathbf{u}_i^{T+1} = -\frac{\partial E}{\partial \mathbf{u}_i}(W^T, U^T) = -\frac{\partial E_i}{\partial \mathbf{u}_i}(W^T, \mathbf{u}_i^T)\right).$$

Update $\mathbf{u}_i$ as follows:

$$\Delta\mathbf{u}_i^{T+1} = (1 - \eta_u) \cdot \Delta\mathbf{u}_i^T + \eta_u \cdot \delta\mathbf{u}_i^{T+1}, \tag{13}$$

$$\mathbf{u}_i^{T+1} = \mathbf{u}_i^T + \alpha_u \cdot \Delta\mathbf{u}_i^{T+1}, \tag{14}$$

where α_u and η_u are learning coefficient and momentum, respectively.

end for

Sum up the delta errors of W for all time steps t of all the time series j of all trained interaction $i \in \mathcal{I}'$ to obtain δW^{T+1}.

$$\left(\because \delta W^{T+1} = -\frac{\partial E}{\partial W}(W^T, U^T)\right).$$

Update W as follows

$$\Delta W^{T+1} = (1 - \eta_w) \cdot \Delta W^T + \eta_w \cdot \delta W^{T+1}, \tag{15}$$

$$W^{T+1} = W^T + \alpha_w \cdot \Delta W^{T+1}, \tag{16}$$

where α_w and η_w are learning coefficient and momentum, respectively.

Testing Recognition and Imitation

The goal of the sNNPB is to re-produce – or imitate – particular object interactions including unfamiliar ones. To test this capability, we simulate a recognition and an imitation process for all the 36 possible interactions $i \in \mathcal{I}$, including unfamiliar ones. In a recognition phase, the sNNPB is presented with 12 randomly generated examples of a particular sensorimotor interaction i. These are conceptualized into a corresponding concept vector $\mathbf{u}_i$ by means of error back-propagation, as used during learning (Alg. 1), where the averaged error determines the recognition vector. Connection weights W are not updated in this phase. Essentially this method induces the re-usage of existing behavioral (sub-)routines rather than the modification of them.

The recognition vector $\mathbf{u}_i$ is then applied in 280 randomly chosen robot-object interaction scenarios, testing its capability to imitate the specific object interaction i. The success rates of all interactions were recorded for later analysis. The successful imitation is judged in terms of the final relative distance and direction to a given target. The robot is required to keep the designated conditions for 100 consecutive time steps within 250 time steps. In order to accomplish the `moveto` interactions, the robot had to stay within 40.0 cm from the target without touching it. For the

`orientto` interactions, the robot was required to orient itself toward the designated direction within an error of $\pm 6^{\circ}$.

Three Network Training Setups

We examine three different variations of the sNNPB. For the basic network, the visual input is encoded in a 27-dimension vector. The visual field is vertically segmented into nine regions. Each region is represented by the fraction of the region covered by colored patches and the dominant hue of the patches in the region. The hue H is encoded by the position $(\cos H, \sin H)$ in the color circle shown in Fig. 4. Motor output is encoded as a two-dimensional vector representing the velocity of the left and right wheels, each of which is a real value ranging from -1.0 to 1.0. A negative value indicates reverse rotation.

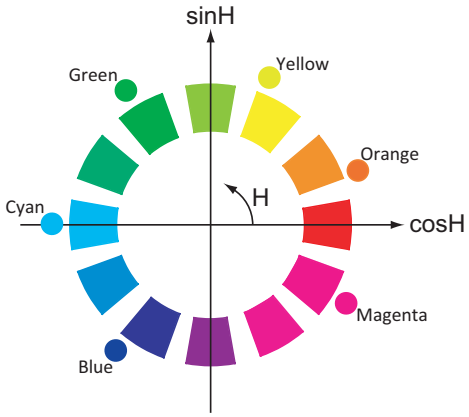

Figure 4. **Colors of the objects**

The two networks other than the basic one have a modified VI layer and MO layer, respectively. The VI-modified network has 36 VI nodes rather than 27, because the dominant color of each region is represented in RGB format. The MO-modified network has 10 MO nodes, instead of 2, encoding the velocity of each wheel by a 5-dimensional vector. This vector represents the velocity v in the form of:

$$G(v) = [f(-0.6; v), f(-0.3; v), f(0.0; v), f(0.3; v), f(0.6; v)],$$

where $f(x; v)$ is a Gaussian distribution $N(x; \mu = v, \sigma^2 = 0.3^2)$.

In both modified cases, the sensorimotor encoding is different from the basic case but *local* similarity is preserved. Motor similarity as defined by the Gaussian and linear encodings are locally identical, since Euclidean distance between Gaussian-encoded velocities $\|G(v_2) - G(v_1)\|$ is almost proportional to $\|v_1 - v_2\|$, given $\|v_1 - v_2\|$ is sufficiently small. For the same reason, sensory similarity defined by hue and RGB encodings are locally identical.

Behavioral Generalizations

This section presents results regarding the obtained behavioral imitation performance. We focus on the generalized imitation capabilities achieved, depending on the distribution of trained

interactions (cf. Fig. 2) and given sensory and motor encodings (see Sec.). We also analyze the generated behavioral mappings over the possible sensory input ranges.

Skill Generalization and Recombination

In the following section we progressively increase the number of trained interactions and analyze the compositionality of the behavioral generalizations achieved. For each block (cf. Fig. 2), three sNNPBs were trained with different network connectivity initializations. Untrained interactions were then tested by first deriving a maximally suitable PB activity by means of back-propagation and then testing that PB activity on 280 other object setups, as detailed in Sec. .

Learning from one interaction only

In Block 1 (cf. Fig. 2), the robot was trained only on `orientto-cyan+0` interaction examples. Figure 5 shows that the acquired skill is, to some extent, transferred across offsets and targets for all three setups. However, no transfer across interaction types (that is, from `orientto` to `moveto`) is realized. The closer the color of the target and the offset of the interaction are, the higher the imitation success rate. The strength of the transfer, however, depends on the given setup. In the RGB-based color encoding case, the transfer across targets is slightly weaker, but it extends a bit further to other color encodings, somewhat reflecting the neighborhood relations of the RGB color space (Fig. 5(a) versus Fig. 5(b)). Similarly, transfer across offsets is more limited in the Gaussian motor encoding case (Fig. 5(a) versus Fig. 5(c)). It appears that the non-linearity induced by the Gaussian motor encoding makes it harder to identify common skills, such as convergent dynamics, among the `orientto` interactions with respect to different offsets.

Note that in almost all of the results the generalization patterns are not perfectly symmetrical. This is mainly due to the randomness induced in the recognition stage, where 12 scenarios are selected randomly independently for each of the 36 interaction types. Thus, while the results point-out obvious tendencies, inevitably a significant amount of noise is inherent.

In Block 2, the `moveto-cyan` interaction was taught. In the linear motor encoding cases shown in Figs. 6(a) and (b), transfer across types of interactions, from `moveto` to `orientto`, is observed, as well as additional offset transfer and color transfer. Thus, convergent dynamics, which are acquired by learning `moveto-cyan`, are reused to generate `orientto-cyan` and neighboring `orientto` behaviors. However, the successful transfer depends on the chosen motor encoding, as no systematic transfer is observable in the Gaussian motor encoding case (c.f. Fig. 6(c)). This shows that while the Gaussian encoding does not diminish learning success, skill-transfer across types is disrupted. Essentially, a successful transfer requires the separation of wheel rotations and forward movement signals, which are linearly combined in the `moveto` interactions (cf. Sec.). This separation is hard to realize in a linear way given the Gaussian motor encoding.

Note that these results indicate that the system detected an *asymmetric, compositional* similarity, since skill-transfer was possible from `moveto` to `orientto` interactions but not vice versa. This asymmetry is due to the behavioral asymmetry between `moveto` and `orientto` interactions: `moveto` requires additional approaching skills, which are not necessary for `orientto`. Both, however, require a pivoting skill, which is blended into the approaching skill in `moveto` interactions. Back-propagation identifies the PB components that control the approaching behavior and thus can inhibit those selectively to realize the untrained `orientto` behaviors.

From a cognitive perspective, this result replicates empirical asymmetry of similarity comparisons (Tversky, 1984) without providing explicit, innate compositional representations. For ex-

(a) Hue + Linear

	blue	cyan	green	yellow	orange	magenta
moveto	00	00	00	00	00	00
orient-30	09	04	24	12	13	12
orient-18	06	45	20	07	11	12
orient+0	33	89	33	07	07	02
orient+18	21	56	53	07	10	08
orient+30	16	17	33	14	10	14

(b) RGB + Linear

	blue	cyan	green	yellow	orange	magenta
moveto	00	00	00	00	00	00
orient-30	15	16	26	15	09	09
orient-18	09	49	23	24	13	09
orient+0	22	99	21	14	06	05
orient+18	30	20	31	31	10	10
orient+30	16	05	29	14	06	07

(c) Hue + Gaussian

	blue	cyan	green	yellow	orange	magenta
moveto	00	00	00	00	00	00
orient-30	02	00	05	09	07	06
orient-18	03	23	09	14	08	03
orient+0	36	91	60	09	06	03
orient+18	44	21	22	03	02	05
orient+30	36	02	08	08	06	09

Figure 5. **Imitation performance obtained in Block 1:** Local offset and color generalizations within the `orientto` interaction space. Success rates of imitation are co-encoded by number and color. Taught interactions are surrounded by a black frame.

ample, we prefer "a scanner is like a copy machine" to "a copy machine is like a scanner." This observation elucidates that similarity in cognition is not always like a typical symmetric distance defined between (pre-defined) representation. Rather, similarity should be defined in terms of how much knowledge, or skill, can be transferred between two things, or tasks.

Compositionality when learning from multiple interactions

Recombinations of interactions are obtained in Block 3, where the `moveto-blue` interaction is trained and all `orientto-0` interactions: the robot can imitate the unfamiliar color-specific `moveto` interactions (Fig. 7(a)). This transfer was not obtained in the Gaussian motor encoding case (Fig. 7(b)), due to the stronger locality of the motor encoding. Nonetheless, it needs to be noted that the color-sensitive orientation component can be diverted from the `orientto` skill into the `moveto` skill, although the `orientto` skill cannot produce the `moveto` skill itself.

This skill transfer can only be explained by a recombination of interactions. It suggests that a sub-symbolic equivalence of the following symbolic compositional system was acquired:

$$\langle Interaction \rangle ::= \oplus(\langle Type \rangle, \langle Target \rangle), \tag{17}$$

$$\langle Type \rangle ::= \texttt{moveto} \mid \texttt{orientto}, \tag{18}$$

$$\langle Target \rangle ::= \texttt{blue} \mid \texttt{cyan} \mid \texttt{green}, \tag{19}$$

where $\langle Interaction \rangle$ is a set of concepts representing possible interactions, and $\langle Type \rangle$ and $\langle Target \rangle$ are a set of elemental concepts representing types of interactions and targets, respectively. A composition rule, which combines a type and a target into a whole interaction, is denoted by $\oplus$.

Acquisition of all interaction skills

Finally, the robot could imitate all the possible interactions well when it learned 21 out of the 36 interactions according to Block 4, as shown in Fig. 8. The generalization capabilities suggest

(a) Hue + Linear

	blue	cyan	green	yellow	orange	magenta
moveto	38	87	35	07	03	09
orient-30	05	02	02	03	04	15
orient-18	14	18	04	02	02	07
orient+0	35	80	14	02	00	00
orient+18	05	15	13	01	01	01
orient+30	21	30	07	01	01	03

(b) RGB + Linear

	blue	cyan	green	yellow	orange	magenta
moveto	42	89	33	00	00	18
orient-30	39	20	05	03	04	08
orient-18	41	64	18	03	01	07
orient+0	39	85	16	12	00	05
orient+18	04	11	21	06	04	03
orient+30	09	03	14	05	12	05

(c) Hue + Gaussian

	blue	cyan	green	yellow	orange	magenta
moveto	29	90	31	04	01	00
orient-30	24	25	05	00	00	01
orient-18	14	11	13	01	01	01
orient+0	02	03	19	03	03	01
orient+18	01	00	08	03	02	01
orient+30	01	00	01	01	02	01

Figure 6. **Imitation performance obtained in Block 2:** Generalizations from `moveto` to `orientto` interactions.

(a) Hue + Linear

	blue	cyan	green	yellow	orange	magenta
moveto	95	91	84	82	81	85
orient-30	68	41	58	61	57	64
orient-18	82	75	63	48	78	78
orient+0	96	91	84	87	87	89
orient+18	64	66	58	70	65	78
orient+30	49	32	60	53	48	61

(b) Hue + Gaussian

	blue	cyan	green	yellow	orange	magenta
moveto	94	53	35	18	11	25
orient-30	40	15	11	22	20	17
orient-18	59	40	16	33	50	27
orient+0	96	88	90	92	87	90
orient+18	43	30	25	39	16	20
orient+30	16	15	17	22	16	17

Figure 7. **Imitation performance obtained in Block 3:** Color-respective generalizations from `orientto` to `moveto` interactions.

that the network developed the following fully compositional system of concepts:

$$\begin{aligned}
\langle Interaction \rangle &::= \oplus_1(\texttt{moveto}, \langle Target \rangle) \\
&\quad | \quad \oplus_2(\texttt{orientto}, \langle Target \rangle, \langle Offset \rangle), \\
\langle Target \rangle &::= \texttt{blue} \mid \texttt{cyan} \mid \texttt{green} \mid \texttt{yellow} \mid \texttt{orange} \mid \texttt{magenta}, \\
\langle Offset \rangle &::= \texttt{-30} \mid \texttt{-18} \mid \texttt{0} \mid \texttt{+18} \mid \texttt{+30},
\end{aligned} \tag{20}$$

where $\langle Interaction \rangle$ is the set of available interaction concepts, and $\langle Target \rangle$ and $\langle Offset \rangle$ are sets of elemental concepts representing target and offset, respectively. Composition rules for the two types of interactions, `moveto` and `orientto`, are denoted by $\oplus_1$ and $\oplus_2$, respectively. The sub-symbolic implementation of this concept system in the sNNPB is analyzed in detail in Section .

Generalized Behavioral Patterns

While the previous section showed that the network is able to generalize its interaction skills in a functionally compositional way, we now focus on how the necessary interaction behavior is generated and how even untrained interactions could be generated. We are interested in the extent to

(a) Hue + Linear

	blue	cyan	green	yellow	orange	magenta
moveto	98	97	98	97	90	85
orient-30	78	77	76	80	80	76
orient-18	76	81	86	76	86	75
orient+0	74	88	92	87	88	74
orient+18	89	76	88	89	74	89
orient+30	73	75	83	85	75	89

(b) Hue + Gaussian

	blue	cyan	green	yellow	orange	magenta
moveto	95	96	95	93	87	88
orient-30	91	83	82	85	89	91
orient-18	92	93	94	84	93	89
orient+0	89	96	97	96	95	93
orient+18	96	94	95	95	87	95
orient+30	82	93	91	93	81	93

Figure 8. **Imitation performance obtained in Block 4:** Both encodings yield a generalization performance in all untrained cells above 73%.

which the sNNPB generalizes the sensorimotor mappings beyond the cases on which it was trained. Therefore, we focus on the replication of and generalization over the object interactions generated by the training program. The presented data in this section were generated based on Block 4 by using an sNNPB with the hue-based, linear encoding setup.

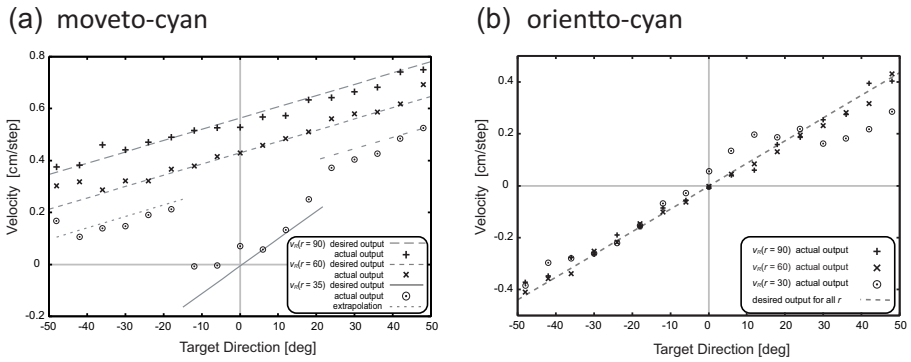

Figure 9. Task-specific direction and distance to velocity mappings.

Figures 9(a,b) show the acquired mappings between the relative direction and distance of a target object and the velocity of the right wheel for `moveto-cyan` and `orientto-cyan`, respectively. The motor output is determined in environments where only one `cyan` object was located. The figures show that the robot can reconstruct the position-velocity mapping for both the `moveto-cyan` and the `orientto-cyan` interactions. An interesting generalization can be observed for the `moveto-cyan` case when the robot is located close to the object ($r = 35$). In this case, the robot was only trained on cases where the object was within 15° relative directional range due to the environmental setup. While the desired output is well-approximated within that range, outside of that range the system extrapolates to the cases where the object is at a further distance. The actual velocity then follows a line labeled "extrapolation" in Fig. 9(a), which is obtained by applying $r = 35$ to the equations for $r \geq 36$ (cf. (3) and (4)).

In addition to this generalization to visual input ranges that were not seen during training, Fig. 10 shows that sNNPB also learned to *avoid* currently undesirable objects. Figure 10(a) shows the hue dependency of motor output v_R for the `orientto-cyan` interaction concept[1]. The output

[1]The motor output v_L was generally mirrored to the one of v_R in this case (not shown).

was recorded by presenting 40 different colored objects at 19 different positions, all 90cm from the robot. At the hue corresponding to the cyan object (hue $\approx 180°$), it can be seen that the robot's right wheel rotates forward when a cyan object is located to the right of the robot but backwards when the object is located to the left, thus replicating the trained sensorimotor interactions. However, if a different colored object is observed given the `orientto-cyan` instruction, this pattern is increasingly reversed. The robot increasingly avoids other objects the stronger the color differs from cyan. These results show that the robot generalized over the color space.

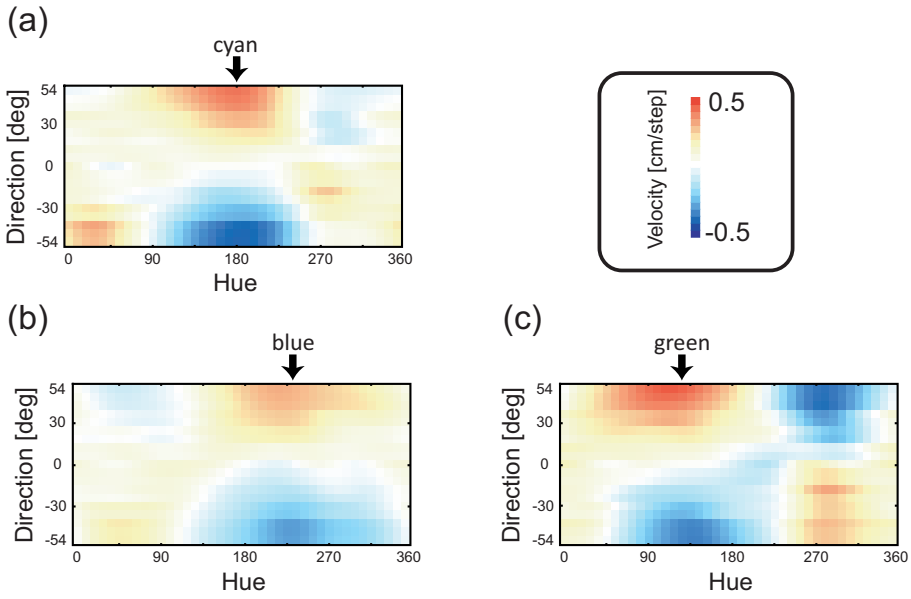

Figure 10. Hue-direction to wheel velocity mappings for specific `orientto` interactions.

Similar patterns can also be observed for other target colors (figures 10(b,c)). It should be noted, however, that `orientto-blue` was never taught during the learning phase. Further analyses revealed that all six `orientto-`τ interactions share a highly similar mapping with different displacements corresponding to the color of the target object τ. This indicates that some common internal mechanisms are reused among all the `orientto-`τ interactions and are modified based on the currently desired target τ, as specified in the PB layer.

Another repetitive motor pattern should be found with respect to the directional offset ϕ, which modified the object-respective turning behavior according to (5) and (6). In Fig. 11, the five lines show the velocity profiles computed from (5) with the trained offsets. The actual direction-velocity mappings for `orientto-cyan` with offsets confirm the existence of the common mechanism. The mappings are obtained under the condition that the `cyan` object is placed 90cm away from the robot. Note that the two interaction concepts `orientto-cyan-30` and `orientto-cyan+18` were never taught during the learning phase. Each of the mappings follows its desired profile. Systematic error patterns can also be observed from the profiles, again suggesting common mechanisms that realize the offset-respective mappings.

In sum, we were able to show (a) behavioral generalizations over the hue space, (b) emergent avoidance behavior that was only indirectly taught, and (c) behavioral generalizations in terms of

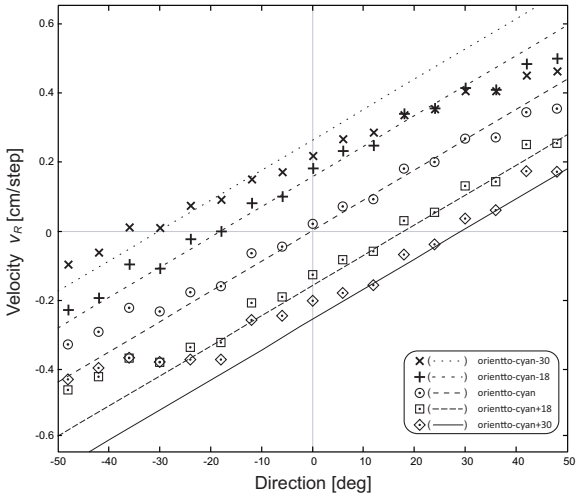

Figure 11. Target direction to wheel velocity map for different offset angles

offset value inter and extrapolations.

Network Structure Analysis

While the behavioral analyses in the last section confirmed that the sNNPB structured its sensorimotor interaction skills in a compositional way, it remained unclear how it accomplished this. It also remained unclear how the network generated the sensorimotor interaction patterns in the first place. To explain this, we now analyze the sNNPB encodings in detail. First, we analyze the structure of the PB layer and show that a sub-symbolic equivalence of compositional role-argument structures emerged. Next, we reveal *how* these compositional structures in the PB layer realize the selective invocation of those sensory-to-motor mappings that realize a particular interaction.

Functional Compositionality in Concept Space

Seeing that the sNNPB can yield behavioral recombination, we first turn to the PB layer and investigate how the interactions are represented in a "compositional" manner. Thus, we structurally analyze the PB space, as defined by the interaction-respective concept vectors $\mathbf{u}_i$ for all the 36 possible interactions $i \in \mathcal{I}$. The presented structures were generated by an sNNPB in the hue-based, linear encoding setup. Comparable structures were obtained by using the other encoding setups.

Figure 12 shows concept structures obtained in Block 1. The 15 concept vectors shown represent `orientto` interactions toward either `blue`, `cyan`, or `green` objects with five different possible offsets. The vectors are projected into a surface spanned by the first two principal components of the 30 concept vectors for the `orientto-`τ`-`ϕ interactions. The six `moveto-`τ concept vectors were removed from this principal component analysis because they took aberrant values. The accumulated contribution rate of the two principal components was 0.77. With regard to the presented 15 interactions, similar interactions are arranged nearby in the concept space and a rudimentary continuum of interactions by offset is observed in a horseshoe shape. However, no regular sub-arrangements are observed, suggesting that the information of target color and offset angle was not clearly sep-

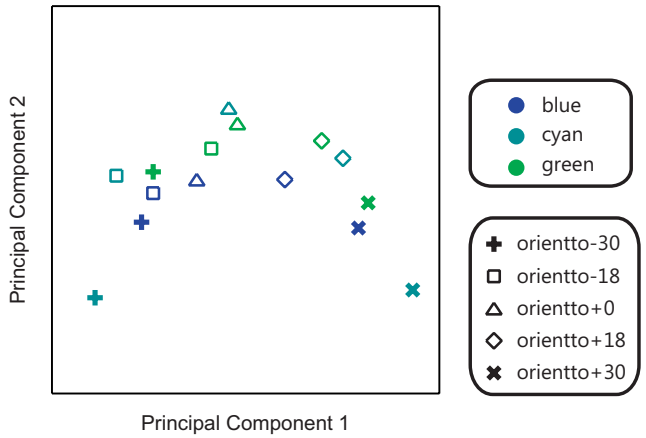

Figure 12. **Concept structure obtained in Block 1:** Mainly sensorimotor similarities are observable.

arated by the sNNPB. This implies that every `orientto` interaction with different offsets employs its proprietary target representation. This is consistent with the observed behavioral performance in the last section: no recombination of interactions was observed.

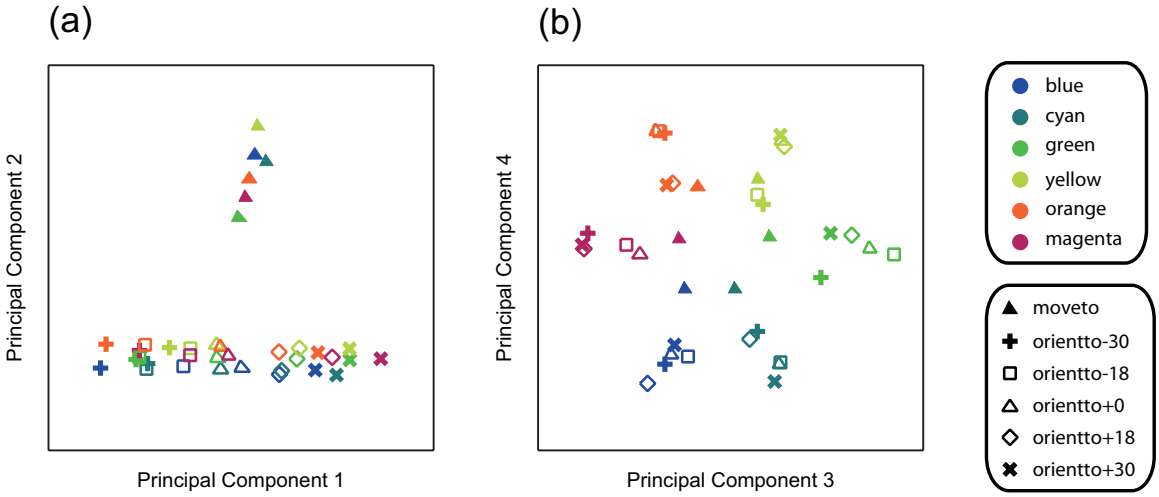

Figure 13. **Concept structure obtained in Block 4:** Full semantic compositionality is observable.

In contrast, in Block 4 a systematic geometric arrangement self-organized among all 36 concept vectors, as shown in Fig. 13. The top four principal components of all 36 concept vectors are presented in the figure, arranged by means of an affine transformation for visualization convenience. The accumulated contribution rate of the first four principle components was 0.78. The arrangement consists of three sub-arrangements, which correspond to the three roles constituting the interactions: $\langle Type \rangle$, $\langle Target \rangle$, and $\langle Offset \rangle$ (cf. Equation 20). In Fig. 13(a), the interactions are clustered with respect to its type, with `moveto` or `orientto` along the y-axis. In the `orientto`

cluster, five sub-clusters can be found for each of the offsets, although they overlap to some extent. These five sub-clusters constitute a linear continuum of `orientto` interactions by offset along the x-axis. In Fig. 13(b), six clusters correspond to each of the target colors. Furthermore, the clusters are arranged in a circle comparable to the circular continuum of color in hue space (see Fig. 4). Further analyses have shown that intermediate arrangements can be found when training other subsets of interactions. Thus, a continuous process underlies the transition from similarity to compositionality, realized by progressively more self-organized geometric arrangements of the concept vectors.

Since the three role-relevant sub-spaces are orthogonal to each other, elemental concepts are reusable in all possible combinations. We discovered the analog mechanism that realizes the role-argument structure estimated in Sec. and specified in (20), considering the following correspondences:

1. Roles $\langle Type \rangle$, $\langle Target \rangle$, and $\langle Offset \rangle$ correspond to respective sub-manifolds.

2. Elemental concepts for each role, for example, `orientto`, `blue`, and `-30`, are vectors pointing to the center of gravity of the corresponding clusters;

3. The argument structure, which combines the elemental concepts, is implemented by the disjoint union of the corresponding vectors.

Thus, roles are sub-manifolds that are embedded in a self-organized, higher-dimensional compositional interaction manifold. Elemental concepts for each role are particular vectors in the corresponding sub-manifold. Compositionality is enabled by activating particular role combinations in the form of disjoint unions, where the orthogonality of the role sub-spaces enables the flexible recombination of actual elemental concepts, which play particular roles in a concrete interaction.

Compositional Sensorimotor Mapping

We have shown that the PB space expresses compositional interaction concepts by representing $\langle Type \rangle$, $\langle Target \rangle$, and $\langle Offset \rangle$ in role-specific subspaces. Since the interaction concept vector determines the mapping from VR to MO (cf. Fig. 3), we now investigate *how* these compositional vectors may activate the appropriate sensorimotor mapping. The crucial questions are: (a) how can the connection weights $\mathrm{sC}_{i,j}$ suitably select the appropriate sensory-to-motor mappings and (b) how is the sensory information transformed in layer VR to enable the weight-driven selection mechanism.

We will show that the compositional behavioral flexibility is realized by the multiplicative re-combination of color-selective sinusoidals stemming from the *sem-net* and the *pb-net*. Activity of the *sem-net* is determined by the currently observed objects and their relative positions. Activity of the *pb-net* is determined by the implicitly given target, interaction type, and angular offset value, which specify a particular interaction. To be more precise in the following mathematical formulation, we denote particular interaction concept weights by $\mathrm{sC}_{i,j}^{\mu\tau\phi}$ with $\mu\tau\phi \in \mathcal{I}$ and interaction type $\mu \in \{\texttt{orientto}, \texttt{moveto}\}$, target $\tau \in \{\texttt{orange}, \texttt{yellow}, \texttt{green}, \texttt{cyan}, \texttt{blue}, \texttt{magenta}\}$, and offset $\phi \in \{-30, -18, 0, 18, 30\}$.

As indicated by the behavioral generalization capabilities and exemplary goal-specific sensor-to-motor mappings depicted in Fig. 10, the sNNPB developed a selectivity to a specific target object τ by rotating a common visuomotor profile along the hue axis, dependent on the hue of the target object ψ_τ. Mathematically, such a rotation can be realized by a sinusoidal function that depends on the target color ψ_τ (in angular hue space) and on the currently observed color ψ_ω (also in angular hue space). This sinusoidal pattern decreases to zero the closer the respective target is situated to the interaction offset angle ϕ. Because turning in opposite directions requires wheel speeds with

opposite signs, a linear dependency on the target angle suffices. Color selectivity and offset angle dependency of the rotation behavior can thus be realized by approximating the following idealized equation:

$$f_{\mu,\tau,\phi}(\psi_\omega, \rho_\omega) = \phi + \rho_\omega \cdot \cos(\psi_\tau - \psi_\omega), \tag{21}$$

where ρ_ω is the angle to an observed object. The cosine part realizes the attraction to a target color and the repulsion to the opposite color. A linear mapping of the color-respective sinusoidal activity to the two wheel motors (with opposite signs) cause the orientation behavior. The offset ϕ can be added to this mapping, effectively shifting the target direction where turning ceases.

To separate the observation-based parts from the task-dependent parts, (21) can be decomposed into two sinusoidals – one of which being 90° phase-shifted from the other. The following formula yields the same result as the one above due to fact that $\cos(x-y) = \cos(x)\cos(y) + \sin(x)sin(y)$:

$$\phi + \cos(\psi_\tau) \cdot \underline{\rho_\omega \cdot \cos(\psi_\omega)} + \sin(\psi_\tau) \cdot \underline{\rho_\omega \cdot \sin(\psi_\omega)}. \tag{22}$$

Now, the static, goal-oriented PB-based selection signals (ψ_τ and ϕ) are separated from the dynamic sensory-based signals, which are underlined in the above equation. The constant orientation offset suggests that the PB-layer may induce an offset-dependent constant activation modulation. The formula also suggests that a directional signal ρ_ω needs to be supplied by the VR nodes, which needs to be blended into the sinusoidal, vision-based activities. This necessitates a directional, color-selective sensor activity, which can result in a Braitenberg-like, but selective, approaching behavior.

Taking a hue-velocity mapping for `orientto-cyan` (Fig. 10a) and its 90° phase shifted mapping as basis mappings, any mappings for `orientto-`τ can be generated based on (22). Note that hue-velocity mappings for `orientto-cyan` draw a sinusoidal curve with respect to ψ_ω for any object direction other than the targeted direction. This sinusoidal pattern does not depend on a particular visual format, since similar results were obtained both for hue and RGB formats. The sinusoidal-determined basis mappings are finally combined with the ψ_τ and ϕ components of the second-order connections, yielding the color-specific wheel velocity mappings depicted in Fig. 10.

In the following two sections we verify that the weights generated by the *pb-net* in layer sC and the processed visual activity generated by the *sem-net* in layer VR exhibit the expected structure for `orientto` behaviors. Signals and further weight biases that are responsible for the approaching component in the `moveto` interactions are also identified.

Selecting and Modifying Interaction Patterns in sC

Equation (22) suggests that the connection weights $\mathrm{sC}_{i,j}$ realize the target selectivity by a sinusoidal distribution pattern over the target color space plus an additive component that specifies the required offset. Additional weight biases for the approaching behavior in `moveto` interactions are to be expected.

Considering the color-respective hue values as input, we approximate the six color-specific weight values along the hue axis with a sinusoidal equation. That is, given a particular interaction type μ, we fit the weight values $\mathrm{sC}_{i,j}^{\mu\tau\phi}$ by the following function:

$$\mathrm{sC}_{i,j} = a_{ij} \cdot \sin(\psi_\tau - b_{ij}) + c_{ij} \;\; \texttt{with} \;\; c_{ij} = A_{ij}\phi + B_{ij}, \tag{23}$$

where ψ_τ denotes the hue of the potential targets τ. Amplitude a_{ij} and phase offset b_{ij} determine the extent and the position of the full sinusoidal period in the hue input space. The fitting curves are obtained by using the conventional Levenberg-Marquardt method.

Figure 14 shows the fitted amplitudes a_{ij}, phase offsets b_{ij}, and function offsets c_{ij} respective the six interaction concepts for all 12 nodes ($j = \{0, 1, \cdots, 11\}$) of the VR layer connecting to the right-wheel neuron ($i = 1$) of the MO layer. In the case of the `orientto` interaction, the corresponding weight values for the left-wheel neuron exhibited highly similar patterns and are not shown. However, in the case of the `moveto` interaction the fits differed significantly for the right- and left-wheel motor neurons, as shown. In addition, Figure 15 shows exemplar sinusoidal mappings as well as the node-respective multiple correlation values of the respective mappings. From the following weight connection analysis we derive node responsibilities in the VR layer.

The weights connecting nodes #1 and #8 with the MO layer show low correlation values and consistently very low amplitude values. Due to these values, these two nodes do not fit into the proposed scheme. We show below that these nodes indeed provide nearly target-color-independent object distance signals. Nodes #3 and #7 show almost equally low correlation values and low amplitudes, which however, do strongly differ with respect to the different offsets. The other eight nodes exhibit high amplitudes and correlation values. Moreover, the phase offset values b are very similar in these nodes over all five modifications of the `orientto` action.

The `moveto` actions also employ somewhat similar sinusoidal patterns. The phase offsets b are generally similar to the ones for the `orientto` interactions, suggesting a similar approach to combining VR activities for the motor output generation. The amplitude values are generally smaller than the values for the `orientto` interactions due to the less aggressive rotation in the `moveto` interaction. However, forward movements need to be superimposed. Only Node #3 shows a higher amplitude compared to the `moveto` interaction values, suggesting a special relevance for `moveto` interactions. The function offsets (Fig. 14(c)) indicate the strong general relevance of nodes #1, #3, and #8 for the `moveto` interactions, compared to the `orientto` interactions. Also note that for these three nodes only, the offset is similar for the mapping to the right wheel and the left wheel, corroborating evidence that these three nodes control the forward motion.

At the same time, the VR nodes are categorized differently in terms of their relevance to the determination of final angular offsets. Strong correlation is observed between the above-mentioned fitting parameter c and a given offset angle ϕ for six VR nodes. Figures 16(a – c) show linear regression results of c in terms of the fitting function $A\phi + B$ where A and B are fitting parameters defined in (23). Examples of fitting curves are presented in Fig. 16(a) for two offset-relevant nodes #2 and #9 and an offset-irrelevant node #10. The categorization is obtained in terms of the proportional constant A and multiple correlation values, as shown in Figs. 16(b) and (c), respectively. The five strongly offset-relevant nodes are fitted with proportional constant ($|A| > 0.002$) and with multiple correlation values higher than 0.98. The offset values of the sinusoidal equation also illustrate this correlation (cf. Fig. 14(c)). It may come as a surprise that the gradients stemming from nodes #2 and #7 are positive while the gradients of nodes #6, #9, and #11 are negative. This suggests that inverted visual information should be contained in the respective VR nodes.

Reusable Components for the Sensorimotor Mapping in VR

The above analysis has shown that the sNNPB utilizes a sinusoidal selection mechanism that determines the relevance of particular VR nodes for particular color-object interactions. An additional value offsets the sinusoidal curves to scale their mapping influence dependent on the desired interaction offset. According to (22), a similar sinusoidal pattern should also be observable in the VR nodes to realize the unrolling of the desired sensorimotor interaction patterns given the sensory input. Originally, we expected that this layer would provide visually-processed factual information

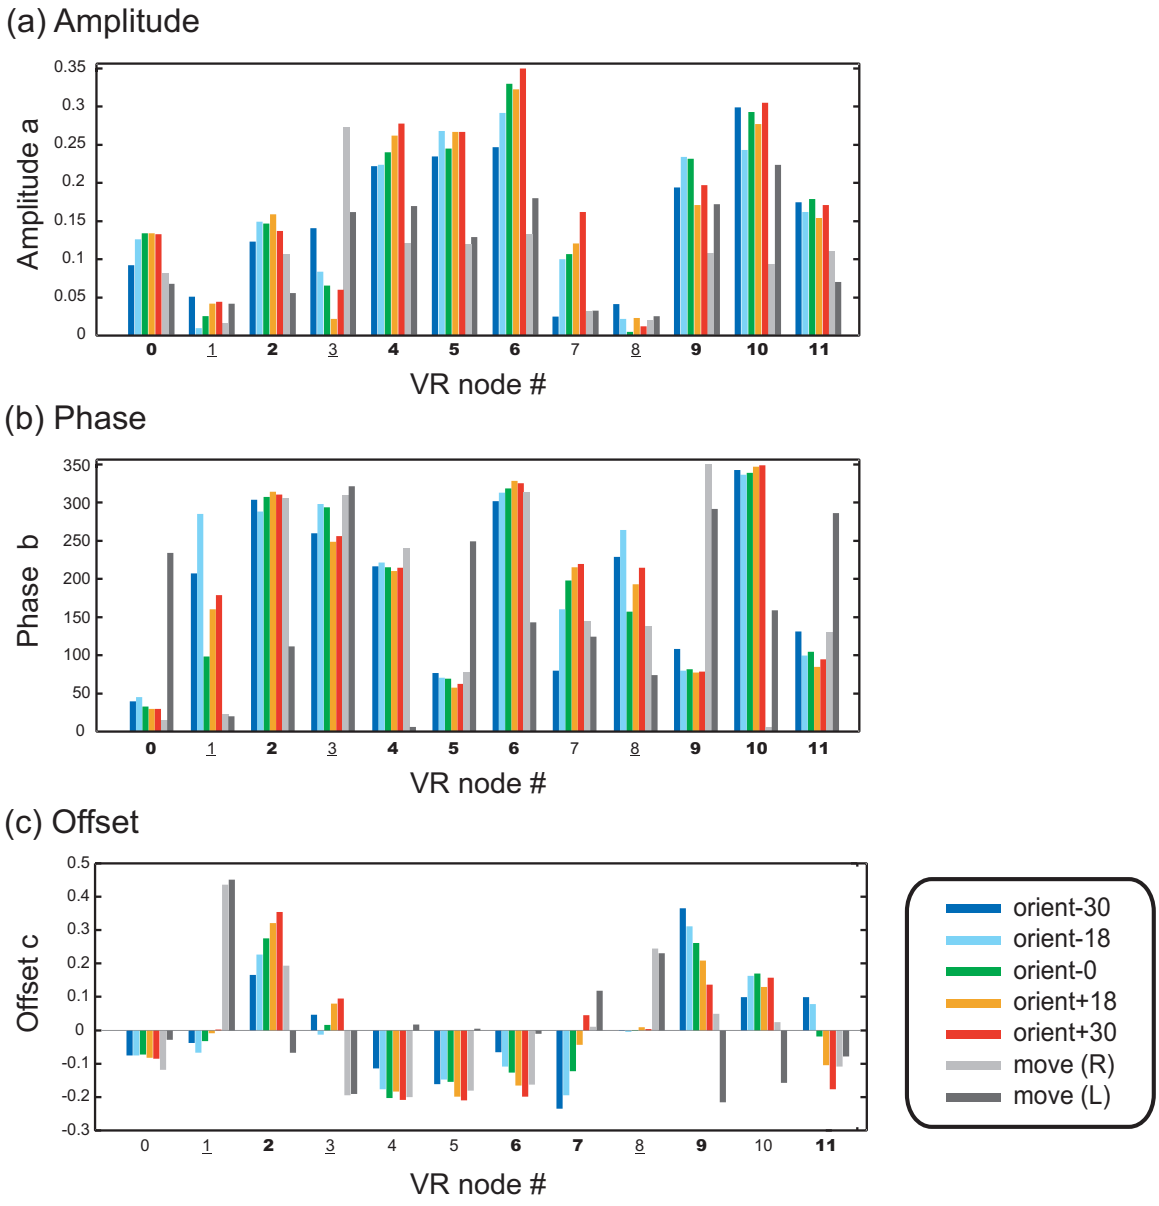

Figure 14. Amplitude (a), phase offset (b), and function offset (c) of each VR node for the six interaction concepts (five `orientto` interactions and the `moveto` interaction): Fitted parameters for the connections to the right-wheel neuron and both right- and left-wheel neurons are presented for the five `orientto` and the `moveto` interactions, respectively. Target-relevant nodes are highlighted in (a) and (b). Offset-relevant nodes are highlighted in (c). `moveto`-specific nodes are underlined.

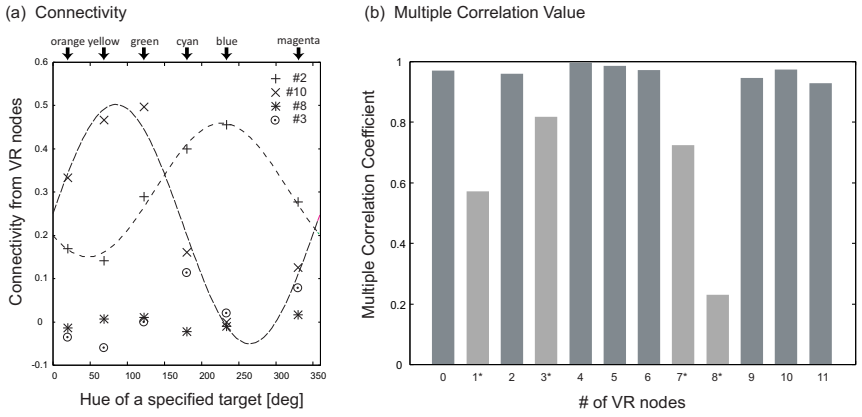

Figure 15. (a) Connectivity from VR nodes to the motor output nodes that generates the velocity of the right wheel; (b) multiple correlation values between actual and approximated connectivity.

about the surrounding environment, such as "orange-object at 45° and blue-object at −30°." However, this layer provides behavior-oriented information about the surrounding objects in the form of processed, Braitenberg-like sensory encodings. These encodings provide the information necessary to realize the underlined parts of Equation (22): sinusoidal patterns along the hue axis are linearly modified dependent on the direction to the object. Distance encodings can also be found that allow the online control of the approaching and stopping behavior necessary for the `moveto` interaction.

We first analyze nodes #1 and #8, which were identified as `moveto`-relevant nodes in the above connectivity analysis. The activation level of #1 decreases gradually as an object comes closer, as shown in Fig. 17. In contrast, the activation of #8 changes suddenly when any object comes just in front of the robot as shown in Fig. 18. This confirms that these two nodes provide suitable distance codes, which can be recruited by the `moveto` interactions. The more negative the activity in these two nodes, the lower the speed of the two wheels. Due to the strong non-linear sensitivity of node #8 around 35cm with a switch from positive to negative values, the sNNPB realizes the stopping-behavior. Essentially, the nodes encode Braitenberg-like distance signals, which allow a direct mapping of object distances onto wheel speeds.

When combining the VR activities of the remaining ten nodes, it can be shown that the sensorimotor mappings for `orientto` interactions are composed of three components. A *common component* is shared by all the `orientto` interactions with any offset angle. An *offset-relevant component* shifts the common component offset-respective. A *target-relevant component* adds the color-dependent sinusoidal tuning. The common component is obtained based on the term B of the sinusoidal function (23). Figure 19(a) shows a sensorimotor mapping that consists of the ten VR nodes with an activity ratio of each node given by the respective Bs. This component is fairly insensitive to hue, but contributes slightly to convergence to any object with no offset.

The offset-relevant component is obtained as a combination of the offset-relevant VR nodes (#2, #6, #7, #9, and #11) with the ratio given by $A\phi$ in (23) and illustrated in Fig. 16. This component is almost flat and contributes to changing the convergent direction by leveling the velocity uniformly as shown in Fig. 19(b) and as was hypothesized in (22). Thus, for each offset-value, the sensorimotor mapping is uniformly shifted to the position where the transfer from positive to

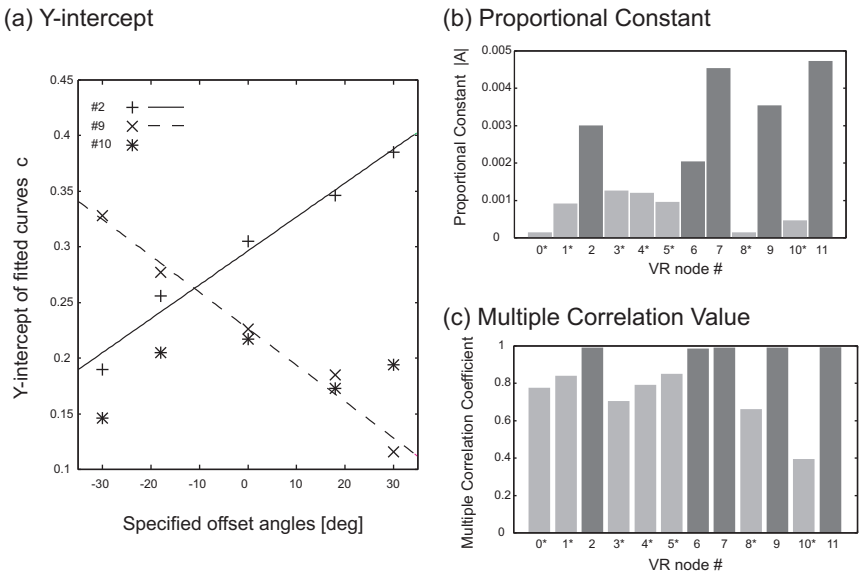

Figure 16. Correlation between c and offset angle

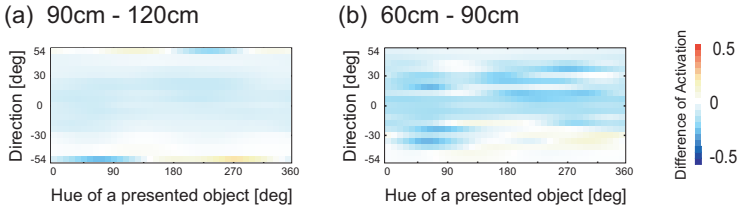

Figure 17. **Activation of** VR **node #1:** Sensitivity of VR node #1 to the distance to an object is shown as differences of two outputs for the object placed at 90cm and 120cm, 60cm and 90cm away from the robot.

negative wheel speeds and vice versa should take place. This encoding suggests the invention of a necessary constant offset value, which can produce constant offsets over the entire sensory space[2].

Finally, the target-relevant component is obtained as a combination of the eight hue-relevant VR nodes (#0, #2, #4, #5, #6, #9, #10, #11) with the ratio given by averaged a and b values over all offset angles. This component produces a suitable linear gradient from positive to negative values. Figures 19(c1 – c3) show that the zero activity hits the spot at the zero offset almost perfectly for the respective target objects. In addition, the pattern increasingly reverses for colors that increasingly differ from the target color. This confirms that the shifting of the sensorimotor mappings observed in Fig. 10 is realized by the multiplicative combination of processed sinusoidal visual information and sinusoidal target selectivity, as formulated in (22). The visual representation encoding can be regarded as the provision of perfectly task-suited, color-sensitive Braitenberg-like sensors: for each target color there are two sensors whose activity is maximal given the target color and decreases

[2]Surprisingly, experiments with bias nodes in the VR layer showed that back-propagation fails to recruit those nodes for the offset control but still "invents" its own offset encoding as explained above.

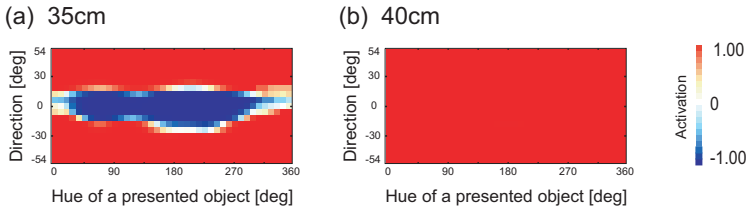

Figure 18. **Activation of** VR **node #8:** the output of VR node #8 is recorded when an object is placed at 35cm and 40cm away from the robot.

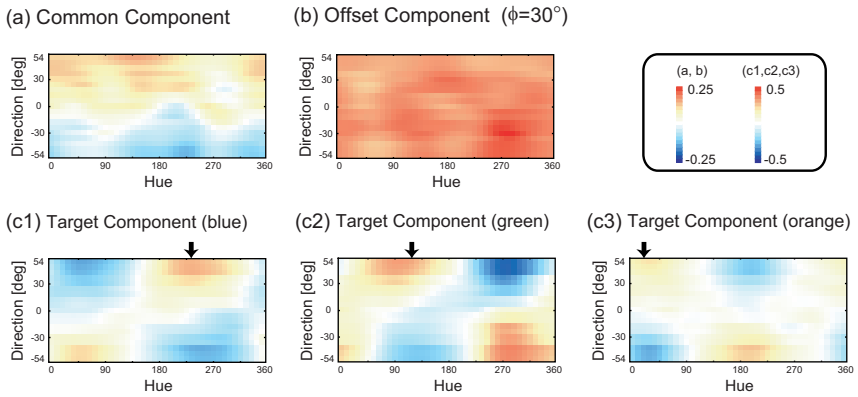

Figure 19. **Reusable Components of sensor-motor maps for `orientto` interactions:** A common component shared by all `orientto` interactions (a), an offset-relevant component (b), and a target-relevant component (c1, c2, c3) are presented.

toward negative values with increasingly different hue-encoded colors. Positioning these sensors with a sufficient perceptual radius to the front-left and -right of the robot can approximately yield the gradient map shown in Figs. 19(c1 – c3).

To see how the visual information is combined to achieve these mappings, we show the output patterns of the other ten nodes of the VR layer in Fig. 20. The output of the VR nodes was recorded by presenting 40 different colored objects at 19 different positions 90cm from the robot. Further analysis reveals that the activation of all ten nodes is approximated by a sinusoidal function of the hue of a presented object ω for particular directions (that is, sinusoidal patterns along the x-Axis for particular target directions). These sinusoidal patterns were observed irrespective of color encodings of the visual input nodes, using either $(\sin\psi_\omega, \cos\psi_\omega)$ (cf. Fig. 20) or (r, g, b) (not shown). Thus, the offset-respective sinusoidal patterns of VR nodes along the hue axis, in conjunction with sinusoidal patterns generated in the $sC_{i,j}$ connections, enable the color selectivity and the observed behavioral generalization over the color space, as mathematically modeled in (22) above.

The linear gradient in (22) comes into being by the PB-controlled combinations of local, color-offset specific linear gradients. These can be observed for given color encodings along the y-Axis in the nodes of Fig. 20, where the observable single non-linearities (switch from positive to negative gradient or vice versa) must be blended to generate an approximate overall linear wheel-

speed mapping (cf. Figs. 19(c1-c3)). Thus, the color-relevant VR nodes provide a mixture of color- and direction-sensitive sensors that do not encode object locations explicitly, but rather encode object directions color-respectively. This encoding allows the necessary selective mapping of directional and distance information onto motor activities, resulting in the observed compositional behavioral competence.

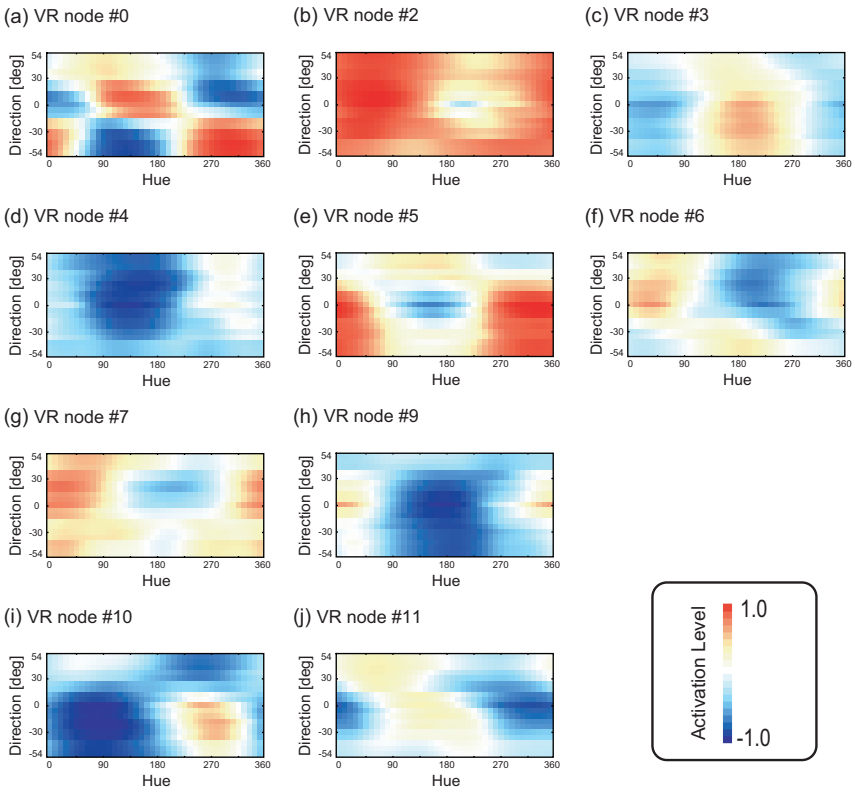

Figure 20. Activation patterns of VR nodes

Developing Compositionality

The previous sections have confirmed that the sNNPB developed compositional concept structures. We now turn to related work to emphasize the novelty of our approach. We focus on the challenges of discovering role-argument structures and symbol-like structures form sensorimotor data alone. Moreover, we discuss the question why the architecture was successful in this endeavor and derive several implications.

Discovering Role-Argument Structures

Previous work has attempted to bridge the gap between analog sensorimotor experiences and a compositional system of discrete concepts. Karmiloff-Smith (1992) proposed implicit *Level-I* representations, which are characterized as implicit, bracketed procedural knowledge. Specific Level-I representations may be acquired by rote-learning from particular sensorimotor events and are thus

analogous to specific experiences. Karmiloff-Smith (1992) then proposes a representational rewriting (RR) model, which converts Level-I representations into increasingly explicit, discrete ones. Because Level-I representations are unstructured, however, they do not provide any mechanism to rewrite themselves into compositional concepts and it remains unclear, how this should be accomplished.

Johnson (1987) and Lakoff (1987) have proposed *image schemas*, which are "derived from our bodily experience" (Johnson, 1987, p.XX). They from compositional concept structures explicitly, because they consist of mutually related concepts derived from those bodily experiences, such as compulsion, containment, or part-whole relations. While they are thus supposed to capture recurring *structural* patterns of our sensorimotor experiences, *how* such concepts are developed autonomously remains unsolved. Clausner (2005) pointed out that a chicken-and-egg problem needs to be solved in this respect: how may image schemas be derived from sensorimotor experiences, while being a prerequisite for organizing these experiences in the first place?

In our experiments, the sNNPB was able to discover role-argument structures by developing role-specific geometric regularities in the form of orthogonal manifold sub-spaces. For example, as abstractly illustrated in Fig. 21, the sub-symbolic equivalence of the `orientto` interaction concept is a cylindrical manifold structure that is particularized by specific target and offset arguments. The target, color-selective role is realized by a ring structure, whereas the offset, orientation-specific role is realized by a perpendicular axis. The perpendicularity enables the independent selection and combination of target and offset. In our experiments, the interaction type formed an additional perpendicular axis, additionally enabling the selective activation of `moveto` or `oriento` interactions (cf. Fig. 13).

These compositional concept structures developed solely based on sensorimotor time series data without a priori schema biases. Sensory and motor similarities were gradually transformed into compositional concept structures, which enabled the combinatorial transfer of skills. With respect to the problem raised in Clausner (2005), we believe that it was actually helpful to refrain from inducing an explicit task decomposition, because this preserved the natural interdependence among semantic roles.

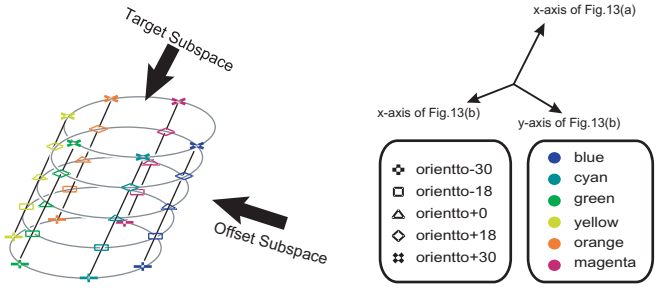

Figure 21. **An idealized concept space:** Three of the four dimensions of concept structure acquired in Block 4 are illustrated.

Structural Emergence in Sub-symbolic Models

Sub-symbolic models that acquire embodied compositional concepts have also been proposed by Cangelosi and Riga (2006) and Tuci, Ferrauto, Massera, and Nolfi (2010). Both models employ

symbolic input layers, which specify interaction concepts that are to be acquired or generated. Each symbolic node corresponds to a particular predefined *elemental* concept. The actual interaction concept is specified by the combination of two or more active nodes. Thus, their models find internal configurations that plug sensorimotor experiences of a robotic agent into predefined concept structures. While our model succeeded without any predefined concept structures, its combination with symbol-oriented representations is pending – potentially contributing another piece to the solution of the symbol grounding problem on a behavioral level (Harnad, 1990; Steels, 2008; Sugita & Butz, 2011).

The previous model in (Sugita & Tani, 2005) actually did utilize linguistic structures to generate a more detailed behavioral conceptualization. In this case, an RNNPB was employed to learn several distinct interactions. The provided linguistic pressure introduced geometric regularities into the behavioral concept space, which made the holistic interaction concepts accessible from the linguistic side. However, no behavioral skill transfer was observed, which was accomplished in the present work.

In the current sNNPB architecture, compositional goal-oriented interaction encodings are fostered because the correlations between target positions and motor outputs form interaction-specific equivalence classes. The sensorimotor similarity alone, however, could not bootstrap the categorization of the interactions, because the sNNPB architecture has no innate mechanism to focus on the target. Neither the visual nor the motor time series data provide similarity signals that may foster the generation of compositional concept structures. Only the provided distinctness information about the trained interactions enabled the network to learn to focus on the target, using the discussed sinusoidal encoding approach.

Pathway Separation and Interaction

Due to the different time scales and the second order connections, the *pb-net* developed goal-oriented interaction selection encodings, while the *sem-net* developed encodings that allowed the activation of any suitable interaction. The second-order weight influence of the *pb-net* onto the sensory-to-motor mapping enforced a multiplicative information combination. In consequence, a compositional encoding developed in the *pb-net*, while directional, goal-differentiating components developed in the *sem-net*. We were not able to achieve similar compositional patterns or behavioral generalizations with previously used RNNPB architectures (Sugita & Tani, 2005). Even continuous-time RNNs (CTRNNs) have not succeeded in discovering similar compositional structures (Yamashita & Tani, 2008; Arie et al., 2010). On the other hand, related work with second-order connections (Nishide et al., 2009) succeeded in generating compositional structures – albeit to a more limited extent. These observations and investigations suggest to us that (a) separating task-specific, goal-oriented encodings from sensory-to-motor encodings and (b) establishing a multiplicative interaction with e.g. second-order connections for the selective activation of goal-suitable sensory-to-motor mappings may be crucial for the development of compositional, cognitive structures.

Interestingly, this separation may be related to encodings that can be found in the dorsal and the ventral sensory processing pathways of the brain (Milner & Goodale, 2008). While the dorsal pathway is characterized as encoding vision-for-action, the ventral pathway is characterized as providing vision-for-representation. In our experiments, re-presentation per se was absolutely irrelevant, but the internal PB structures developed "pro-presentations" for interaction selection. Therefore, we believe that the *sem-net* exhibits properties of the dorsal "perception-for-action" pathway,

whereas the *pb-net* exhibits properties of the (deep) ventral "perception-for-selection of interaction" pathway.

Body-relative, multi-sensory stimulus encodings can be found in the parietal and premotor cortex (Holmes & Spence, 2004; Rizzolatti, Fadiga, Fogassi, & Gallese, 1997), which are transferred into goal-oriented, partially ethologically relevant motor encodings (Graziano, 2006) along the dorsal processing pathway. Along this pathway (but also elsewhere) *gain fields* were identified, which realize behavior-oriented coordinate transformations (Andersen, Snyder, Bradley, & Xing, 1997; Chang, Papadimitriou, & Snyder, 2009; Salinas & Sejnowski, 2001). Gain fields are assumed to multiply receptive field activities from one domain with a monotonic function stemming from another domain (Chang et al., 2009). The *sem-net* in our sNNPB appears to develop similar encodings. In particular, the VR layer developed several neurons that exhibited sinusoidal, receptive-field-like object color activities that were multiplicatively modified by object directions. To accomplish this, the sNNPB converted the provided "retinotopic" sensory inputs into direction signals – encodings found along the dorsal pathway in the primary visual cortex and various parts of the posterior parietal cortex, respectively.

The ventral pathway, on the other hand, is assumed to encode object identities (Riesenhuber & Poggio, 2000) and leads to the selection of currently desired environmental interactions. While we do not implement the object identification part, the *pb-net* mimics the selective activation of a particular, currently desired interaction. While very little is known about the systematicity of the encodings found in the deep ventral pathway, flexible, compositional representations are necessary to effectively coordinate motor interactions. The sNNPB architecture developed such compositional representations in its PB neurons. This suggests that learning biases towards semantically compositional structures may be inherent when separating the processing of sensory information for the control of potential behavioral interactions from representations suitable for the selection of a particular interaction. Consequently, the (deep ventral) action selection pathway may be especially predestined to develop compositional structures.

Summary and Conclusion

In this paper we have shown that a second-order neural network with parametric biases (sNNPB) can develop compositional concept structures for the selective, compositional invocation and control of environmental interactions. This was achieved by training the sNNPB on exemplary sensorimotor time series data using the general back-propagation learning algorithm and providing distinctness (but not compositional) information for each type of trained interaction. The compositional behavioral competence was realized by two multiplicatively interacting modules: the *pb-net* self-organized the neural PB space into a geometrically-structured, compositional concept space, determining the task-relevant sensory-to-motor mapping. The *sem-net* learned to transfer raw sensory information into task-oriented Braitenberg-like pro-motor encodings, providing suitable sensory information for compositional environmental interactions. The multiplicative combination of the two sources of information resulted in the capability of imitating even untrained behavioral interactions, generalizing the learned capabilities in a compositional manner. Since experiments with other neural networks including RNNPBs Sugita and Tani (2005) and CTRNNs (Yamashita & Tani, 2008; Arie et al., 2010) did not succeed in exhibiting similar compositional encodings or behavioral generalizations, but another second-order multiple time-scales RNN did to some extent (Nishide et al., 2009), it appears that multiplicative, second-order connections are indeed necessary to learn compositional structures from sensorimotor interaction patterns.

Future work will need to refine the current approach in several respects. First, it is unsatisfactory that the distinctness information is provided by a teacher. We believe that distinctness information may come either from internal motivation and reward signals or from linguistic signals. The self-organized behavioral PB structure may then serve as a pre-linguistic concept structure, which may facilitate later syntax acquisition (Dominey, 2006, cf.) and behavior-based symbol grounding Harnad (1990).

Second, the model should be extended to acquire recursively structured interaction concepts. In this way, the same target concept may, for example, play the role of subject or object. The recursive structure may be represented as a fractal geometric arrangement, as shown in Nishimoto and Tani (2004) and Sugita and Butz (2010). A learning model in which such a fractal structure self-organizes through the learning of sensorimotor time series involving behaviors of an agent still needs to be developed.

Finally, to foster the scalability of the system to more diverse and complex interactions, it will be necessary to further modularize the learning architecture and to introduce more explicit mechanisms of focus, goal-directedness, and attention. To be able to identify the interaction-crucial bits of information, redundant sensory, motor, and sensorimotor encodings will be necessary. Clearly, a major challenge for future research is to learn compositional structures in more general environmental setups and with a more capable robotic system. The insights gained by the current work suggest that the task-dependent, second-order control of mappings from suitably processed, promotor encoded sensory information to motor activity encodings may play a crucial role in achieving this endeavor.

Acknowledgment

This study is supported by Grants-in-Aid for Scientific Research (No. 20700188) from the Japan Society for the Promotion of Science (JSPS). Yuuya Sugita and Martin Butz also acknowledge funding from the Emmy Noether program (German Research Foundation, DFG, BU1335/3-1) and thank the COBOSLAB team.

References

Andersen, R. A., Snyder, L. H., Bradley, D. C., & Xing, J. (1997). Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual Review of Neuroscience*, *20*, 303-330.

Arie, H., Endo, T., Jeong, S., Lee, M., Sugano, S., & Tani, J. (2010). Integrative learning between language and action: A neuro-robotics experiment. *20th Int. Conf. on Artificial Neural Networks*, *ICANN2010*, 256-265.

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 577–600.

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.

Cangelosi, A., & Riga, T. (2006). An embodied model for sensorimotor grounding and grounding transfer: Experiments with epigenetic robots. *Cognitive Science*, *30*(4), 673-689.

Caruana, R. (1997). Multitask learning. *Machine Learning*, *28*, 41–75.

Chang, S. W. C., Papadimitriou, C., & Snyder, L. H. (2009). Using a compound gain field to compute a reach plan. *Neuron*, *64*, 744 - 755.

Clausner, T. (2005). Image schema paradoxes: Implications for cognitive semantics. In B. Hampe (Ed.), *From perception to meaning: Image schemas in cognitive linguistics* (pp. 93–110). Berlin: Mouton de Gruyter.

Dominey, P. (2006). From holophrases to abstract grammatical constructions: insights from simulation studies. In E. Clark & B. Kelly (Eds.), *Construction in Acquisition* (pp. 137–162). Stanford CA: CSLI Publications.

Fodor, J., & Lepore, E. (2002). *The Compositionality Papers*. Great Clarendon Street, Oxford: Oxford University Press.

Gentner, D., & Markman, A. (1997). Structure Mapping in Analogy and Similarity. *American Psychologist*, *52*(1), 45–56.

Graziano, M. S. A. (2006). The organization of behavioral repertoire in motor cortex. *Annual Review of Neuroscience*, *29*, 105-134.

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, *42*, 335–346.

Holmes, N. P., & Spence, C. (2004). The body schema and multisensory representation(s) of peripersonal space. *Cognitive Processing*, *5*, 94-105.

Johnson, M. (1987). *The body in the mind: The bodily basis of meaning, imagination and reason*. Chicago: University of Chicago Press.

Karmiloff-Smith, A. (1992). *Beyond modularity: A developmental perspective on cognitive science*. Cambridge, MA: MIT Press.

Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.

Mandler, J. M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review*, *99*, 587-604.

Markman, A., & Stilwell, C. (2001). Role-governed categories. *Journal of Experimental and Theoretical Artificial Intelligence*, *13*(4), 329–358.

Meltzoff, A. (1988a). Imitation, objects, tools, and the rudiments of language in human ontogeny. *Human Evolution*, *3*, 45–64.

Meltzoff, A. (1988b). Infant imitation and memory: Nine-month-olds in immediate and deferred tests. *Child Development*, *59*, 217–225.

Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, *46*(3), 774 - 785. (Consciousness and Perception: Insights and Hindsights - A Festschrift in Honour of Larry Weiskrantz)

Nelson, K. (1988). Where do taxonomic categories come from? *Human Development*, *31*, 3-10.

Nishide, S., Nakagawa, T., Ogata, T., Tani, J., Takahashi, T., & Okuno, H. G. (2009). Modeling tool-body assimilation using second-order recurrent neural network. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, *(IROS-2009)*, 5376-5381.

Nishimoto, R., & Tani, J. (2004). Learning to Generate Combinatorial Action Sequences utilizing the Initial Sensitivity of Deterministic Dynamical Systems. *Neural Networks*, *17*(7), 925–933.

Pastra, K., & Aloimonos, Y. (2011). The minimalist grammar of action. *Philosophical Transactions of the Royal Society B: Biological Sciences*. (in press)

Pollack, J. (1990). Recursive Distributed Representations. *Artificial Intelligence*, *46*, 77–105.

Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience*, *3*, 1199-1204.

Rizzolatti, G., Fadiga, L., Fogassi, L., & Gallese, V. (1997). Enhanced: The space around us. *Science*, *277*, 190-191.

Salinas, E., & Sejnowski, T. J. (2001). Gain modulation in the central nervous system: Where behavior, neurophysiology, and computation meet. *The Neuroscientist*, *7*, 430-440.

Steels, L. (2008). The symbol grounding problem has been solved. So what's next? In M. de Vega, A. M. Glenberg, & A. C. Graesser (Eds.), *Symbols and embodiment: Debates on meaning and cognition* (p. 223-244). New Haven: Academic Press.

Sugita, Y., & Butz, M. (2010). *Towards Emergent Strong Systematicity in a Simple Dynamical Connectionitst Network.* (Presented at Workshop on Cognitive and neural models for automated processing of speech and text (CONAS2010))

Sugita, Y., & Butz, M. V. (2011). Compositionality and embodiment in harmony. In P.-Y. Oudeyer (Ed.), *Amd newsletter* (Vol. 8, p. 8-9). IEEE CIS.

Sugita, Y., & Tani, J. (2005). Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adaptive Behavior*, *13*, 33-52.

Tani, J., Ito, M., & Sugita, Y. (2004). Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using RNNPB. *Neural Networks*, *17*, 1273–1289.

Thrun, S., & O'Sullivan, J. (1996). Discovering structure in multiple learning tasks: The TC algorithm. *Proceedings of the Thirteenth International Conference on Machine Learning*, 489–497.

Tuci, E., Ferrauto, T., Massera, G., & Nolfi, S. (2010). Co-development of linguistic and behavioural skills: Compositional semantics and behaviour generalisation. In *Proceedings of the 11th int. conf. on simulation of adaptive behavior (sab2010)* (pp. 523–532). Berlin: Springer-Verlag.

Tversky, A. (1984). Features of Similarity. *Psychological Review*, *84*, 327–352.

Yamashita, Y., & Tani, J. (2008, 11). Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *PLoS Computational Biology*, *4*(11), e1000220.