

Noname manuscript No.
(will be inserted by the editor)

Sensorimotor Input as a Language Generalisation Tool

A Neurorobotics Model for Generation and Generalisation of Noun-Verb Combinations with Sensorimotor Inputs

Junpei Zhong · Martin Peniak ·
Jun Tani · Tetsuya Ogata ·
Angelo Cangelosi

Received: date / Accepted: date

Abstract The paper presents a neurorobotics cognitive model to explain the understanding and generalisation of nouns and verbs combinations when a vocal command consisting of a verb-noun sentence is provided to a humanoid robot. This generalisation process is done via the grounding process: different objects are being interacted, and associated, with different motor behaviours, following a learning approach inspired by developmental language acquisition in infants. This cognitive model is based on Multiple Time-scale Recurrent Neural Networks (MTRNN). With the data obtained from object manipulation tasks with a humanoid robot platform, the robotic agent implemented with this model can ground the primitive embodied structure of verbs through training with verb-noun combination samples. Moreover, we show that a functional hierarchical architecture, based on MTRNN, is able to generalise and produce novel combinations of noun-verb sentences. Further analyses of the learned network dynamics and representations also demonstrate how the generalisation is possible via the exploitation of this functional hierarchical recurrent network.

Keywords Recurrent Artificial Neural Networks · Language Learning · Multiple Time-scale Recurrent Neural Network · Developmental Robotics · Neurorobotics

1 Introduction

For the design of social robots [4,13], besides of building robots with human-like external morphology, the ability to process, to understand and generate language is one of the key factors to support human-robot interaction. However, to build a model to accomplish similar processes for social robotics, the design of the robot’s abilities of understanding, generation and generalisation of natural language is still an open challenge. Particularly, natural language understanding for a social robotic system plays an essential role as it interfaces the vocal command from human users to an internal representation in the robot’s own cognitive system. In this study we will follow a developmental robotics approach to the design of language and communication abilities in robots, following an incremental and interactive process to language learning, inspired by language development in infants.

1.1 Language Understanding for Robot Systems

Important recent developments in social robotics, such as robots performing human-like emotion expression [76] and social attention for autonomous movement [45], have been accompanied by language understanding approaches focusing on the grounding of natural language into the agent’s sensorimotor experience and its situated interaction [5,60].

For instance, in [66,39], syntactic parsing techniques are used to ground the language into primitive motor actions (e.g., pickup, move, place), which can be inferred within graph models. Similarly, Misra et al. [43] developed a system

J. Zhong
Department of Intermedia Art and Science, Waseda University, 3-4-1
Ohkubo, Shinjuku, Tokyo, Japan, 169-8555
Centre for Robotics and Neural Systems, University of Plymouth, PL4
8AA, United Kingdom
Tel.: +44 (0)175284908, E-mail: zhong@junpei.eu
M. Peniak
Cortexica Vision Systems, London, United Kingdom
J. Tani
Korea Advanced Institute of Science and Technology, Daejeon, South
Korea
T. Ogata
Department of Intermedia Art and Science, Waseda University, Tokyo,
Japan
A. Cangelosi
Centre for Robotics and Neural Systems, University of Plymouth,
United Kingdom

for mobile robots which is able to learn to ground the language instructions from a corpus of pairs of natural language including both verbs and spatial information. [74] proposed that in order to understand the object affordance which can be described by adjectives, the most crucial property is the shape-related one.

Besides the direct modelling methods for robot language learning, an alternative approach to build a learning model for language is based on developmental robotics [71, 1, 10]. Taking inspiration from developmental psychology and developmental neuroscience studies, this approach emphasises the role of the environment and of the interactions that occur during learning, over a progression of learning stages. In the context of language understanding, the core of developmental robotics approaches to language learning is following a similar developmental pathway of infants acquiring grounded representations of natural language and forming a symbol system through embodied interaction with the physical environment [6]. Furthermore, via language learning an agent should also be able to generalise by inferring un-trained combinations of words within the lexical constructions acquired.

Various developmental robotics models have been developed that incrementally model the various stages of language acquisition in infants, from phoneme acquisition, to object and action names, to word combinations. For example, the cognitive model presented in [21] outlines the cortical interactions in the syllable generation process which result in different developmental phenomena. This mimics the first stage of language development. The Elija model [28] is a vocal apparatus which strictly follows detailed developmental stages. Working as an articulatory synthesizer, it firstly learns the production of sounds on its own. Then a caregiver is used to produce speech by using speech sounds for object names using reinforcement learning, where the reward is again given by the response of the caregiver. Likewise, a self-organizing map together with reinforcement learning was proposed in [70], which demonstrated that the reinforcement learning based on the similarity of vocalization can improve the post-learning production of the sound of one’s language.

From the models mentioned above, we can see that most of the methods for modelling the first stages of phonetics production do not tend to use robotic platforms. On the other hand, for the modelling of the later stages of lexical development, after assuming that phonetics skills are mastered, robotic systems are usually employed to establish the meta-knowledge about the association between vocal speech and the referents or the actions. Therefore, except studies focusing on the mental imagination of actions as in [20], the mechanical morphology of a robot is particularly important when modelling the acquisition of words, especially those used to name the motor actions. For instance, the model from [38] gets as input dance-like combinations of human

movement primitives plus ambiguous labels associated with these movements. Concentrating on the second and third stages of associating lexicon, words and motor actions, the robot in [16] is able to acquire new motor behaviours in an on-line fashion by grounding the vocal commands on the pre-defined control motor primitives. Similarly, Siskind [57] proposed a model which uses visual primitives to encode notions of different actions to ground the semantics of events for verb learning. Using structured connectionist models (SCMs), [11] built a layered connectionist model to connect embodied representations and simulative inference for verbs. In [8], the emergence of verb-noun separation is learned while the agents are interacting and manipulating the objects.

[61] further developed the grounding the verbs with more complex meanings (such as “keep”, “reject”, “accept” and “give”) which related to the internal states of the caregivers and which were used to build a robotic model for the grounding of increasingly abstract motor concepts and words. As follow-up studies of [16], [15, 25] focused on the understanding of grammatical complexity. They used recurrent neural networks (RNN) to learn grammatical structure based on temporal series learning in artificial neural networks. Also using RNN, [62] reported experiments with a mobile robot implementing a two-level RNN architecture called Recurrent Neural Network with Parametric Bias Units (RNNPB). This allows the robot to map a linguistic command containing verbs and nouns into context-dependent behaviours corresponding to the verb and noun descriptions respectively. Comparing to RNNPB, another kind of RNN architecture called Multiple Timescale Neural Network (MTRNN) is able to ground different scales of sensorimotor information into the hierarchical structure of sentences, such as the spelling of words [46] and words and sentences [26]. The kind of recurrent models provide a memory to store the spatial and temporal structure of the environment and the lexical structures. Given the fact that RNN can learn arbitrary length of the dependencies in statistical structures and their context, the storage ability of the RNN out-performs most of the language learning models.

On a higher level, concerning the meta-learning principles in these learning models, some developmental studies of language have focused on the intrinsic motivation to learn to speak, not only through reinforcement learning, but also following child psychology evidence that language learning can be driven without the explicit rewards from the caregivers. For instance, language commands can be acquired from learning from demonstration (LfD) [55], intrinsic plasticity [48] and evolution [59, 14]. These intrinsic motivation capabilities are implemented through learning models, which allow the agents to acquire communication skills through vocal interactions, besides the use of reinforcement learning techniques, such as the learning by demonstration model for grounding vocal commands into situated spatial informa-

tion [18]; a comparative study for evolving robot language with situated information can be found in [49]. These models demonstrate different forces underlying language learning. From a mathematical perspective, those learning methods overcome the natural language learning bottlenecks of building compositionality of lexical structures and maximising the observed content [58] by means of statistical methods (LfD), optimal control (Reinforcement Learning) or other methodologies.

1.2 Cognitive Background and Motivation

In the developmental psychology studies focusing specifically on the learning of nouns and verbs, there is still an open debate between the learning stages and their relative temporal acquisition order. For the early stages of verb and noun learning, it is widely accepted that most of the common nouns are generally learned before verbs [19], by first connecting speech sounds (labels, nouns) to physical objects in view. However, some nouns which relate to context, such as “passenger”, are learnt relatively late, only after “an extensive range of situations” (contexts or life phases) have been encountered [22], during which verbs may play a crucial role. The embodied learning of verbs and nouns is not correlated to one single modality in sensory percept’s: experiments done in [32] suggest that the nouns are grounded from the intrinsic properties of an object, even at different movements and orientations, while verbs are accounted for the movement path of an object. This distinction may be associated with the neuroanatomy distinction between the ventral and dorsal (what/where) visual streams, involved in the generation of nouns and verbs respectively. As [37] suggested, some nouns and verbs can be learnt more straightforward to learn because they can be accessed perceptually. On the other hand, some abstract words, either verbs or nouns, should only be learnt from a social and linguistic context.

For instance, while infants learn the word-gesture combination at the age of two, they associate the meaning of verbs with the meanings of the higher-order nouns [3]. Such verbs with complex meaning are obtained from both motor action and visual percept [36]. As summarised in [7,9], comparing to the static object perception that associates to simple nouns, the early verb learning involves a temporal dynamic from motion perception. Indeed, we assert that the learning processes of nouns and verbs (especially for those with complex meanings) are not separated; there is a close relation between verb and noun development, during which the embodied sensorimotor information plays a crucial role. During this embodied development, both the perceptual system and the motor system contribute to language comprehension (e.g. [53,31,50,56]). These experiments also extend Piaget’s proposal that language learning is a symbolised understanding process for

dynamic actions, which is “a situated process, function of the content, the context, the activity and the goal of the learner”.

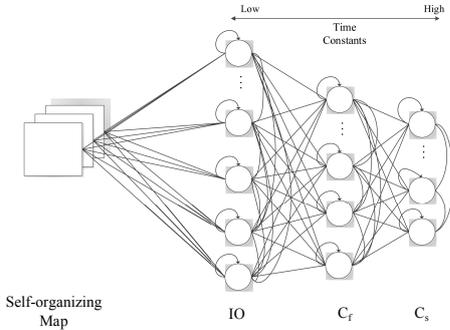
The sensorimotor information is not the only mechanism acting as a learning tool for language acquisition. Conversely, recent research also proposes that language is such a flexible and efficient system for symbolic manipulation which is more than a communication tool of our thoughts (e.g. [35, 41,42].) For the predictive effect from language to sensorimotor behaviours, vocal communication can be one of the sources that drives the visual attention to become predictive, by making inferences as to the source-inferences [67]. In this process, language can trigger a predictive inference about the appearance of a visual percept, driving a predictive saccade [17]. Therefore, the sensorimotor system is affected by the inferences from the auditory modality or even from higher level cognitive processes.

Following the hierarchical cognitive architecture proposed in [75], the language understanding can be represented hierarchically from the neural processes on the (lower) receptor level to the higher level understanding which happens in the (higher) prefrontal cortex. Moreover, we will use a hierarchical recurrent neural architecture, as in [78–80], due to the fact that the learning modalities of visual perception and motor actions can be represented as both spatial and temporal sequences, so that the recurrent connections provide possibilities to intertwine these two modalities. In this paper, the MTRNN model will be employed to ground the features from different modalities with language structures in different time-scales. Similar RNNPB [62] or MTRNN [23] networks have been used to learn verbs and nouns features with motor actions and visual features. The proposed model will use a single MTRNN model to learn both the sensory and motor information in a single set of sequences. We regard the perception and action having inseparable links (e.g. [72, 44]) and should be encoded solely as similar data structures. Moreover, the training of such a large MTRNN has become more and more feasible in recent years due to the accessibility and affordability of GPU computing. Therefore, the two modalities of our MTRNN can be conceptualised at the same time over the embodied sensorimotor experiences towards abstract and compact representations on the higher level of this hierarchy, similarly to the developmental processes of language conceptualisation and categorisation.

2 The Multiple Timescale Recurrent Neural Network Model

This model is based on the combination of a MTRNN network with Self-Organizing Maps (SOMs) to control the humanoid robot iCub, being trained on the understanding of a set of noun-verb combinations to perform a variety of actions with different objects. Fig. 1 shows the learning architecture

Fig. 1: Architecture of Multiple Time-scale Recurrent Neural Network



incorporating a Multiple Timescale Recurrent Neural Network (MTRNN) [73] and the self-organizing maps. The core module of the system is the MTRNN, which will learn sequences of verb-noun instructions and will control the movement of the robot in response to such instructions. The inputs to the MTRNN correspond to the language command inputs, to the visual inputs as well as the proprioceptive inputs. We regard these three modalities as a whole sensorimotor input because the MTRNN model is able to learn the relation between verbs and nouns and seen objects within the context of the non-linearity of the sensorimotor sequences in a hierarchical manner. This network will learn this non-linearity in the functional hierarchy in which the neural activities are self-organised, exploiting the spatiotemporal variations.

2.1 Using a Self-organizing Map as a Sparse Structure

The initial input data sets, consisting of speech, camera images, and kinaesthetic imitation proprioceptive states are pre-processed (See Eqs. 1- 4) using three SOMs respectively for the linguistic, visual and motor input modalities.

Although the MTRNN could be trained with original data representation, we usually employ pre-processing modules for the MTRNN inputs, which results in a sparse structure of the weighting matrices in the network. Also the MTRNN outputs are decoded into the original data structures. The sparseness in weighting matrices has a similar concept of sparse coding in computational neuroscience [47]. The weighting matrices are sparsely distributed, which is an analogous form of the sparse distributed representations that are used in our brain, such as in visual [68] and auditory cortex [54]. Previous research on language learning in RNN [2] also showed that a sparse encoding results in robustness in training and a better generalisation results and improved robustness with noisy inputs.

Here the sparseness structure in the weight matrices is given by the SOMs [34]. During this process, the SOM performs as a dimensional mapping function, with an output

space with higher dimensions than the input space. Having a discretised and distributed neural encoding in the output space, the pre-processed SOM modules are able to reduce the possible overlap of the original data within the original input space. Therefore, the topological homomorphism produced by the SOM guarantees that the training vectors between the raw training-sets and the input vectors are topologically similar with each other.

In the SOM training here, assuming the input vectors are

$$x = [x^1, x^2, \dots, x^m]^T \quad (1)$$

These input vectors are mapped to an output space whose coordinates define the output topology of the SOM. Connecting between the input and output spaces, the weight vector is defined as

$$w_j = [w_j^1, w_j^2, \dots, w_j^m]^T, j = 1, 2, 3, \dots, n \quad (2)$$

where neuron j is one of the input space vectors and n is the total number of those neurons. When a self-organising map receives an input vector, the algorithm finds a neuron associated with weights that are most similar to the input vector. The measure of similarity is usually done using the Euclidean distance metric, which is mathematically equivalent to finding a neuron with the largest inner product $w_j^T x$. Thus the very neuron that is the most similar match for the input vector is referred to as best matching unit (BMU) and it is defined as:

$$c = \arg \min_j \|x - w_j\| \quad (3)$$

The dimensionality mapping is achieved when the BMU coordinates are used to update the weights of the neighbourhood neurons around neuron c by driving them closer to the input vector at iteration t :

$$w_j(t+1) = w_j(t) + \delta(x_j - w_j) \quad (4)$$

where δ is a neighbourhood function from the distance from BMU.

Therefore, the output of the SOM which is encoded in a high-dimensional input space, is still able to preserve the topological properties of the input space due to the use of the neighbourhood function.

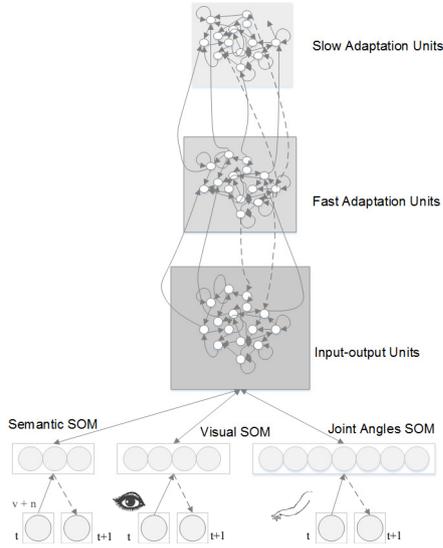
2.2 Multiple Timescale Recurrent Network (MTRNN)

As shown in Fig. 2, the neurons in the MTRNN form three layers: an input-output layer (IO) and two context layers called Context fast (C_f) and Context slow (C_s). In the following text, we denote the indices of these neurons as:

$$I_{all} = I_{IO} \cup I_{C_f} \cup I_{C_s} \quad (5)$$

where I_{IO} represents the indices to the neurons at the input-output layer, I_{C_f} belongs to the neurons at the context fast layer and I_{C_s} belongs to the neurons at the context slow layer. The neurons on a layer own full connectivity to all neurons within the same and adjacent layers, as shown in Fig. 1. The difference between the fast and slow context layers as well as the input-output layer consists in having distinct time constants τ , which determine the speed of the adaptation given a time sequence with a specific length, when updating the neural activity. The larger the value of τ , the slower the neuron adaptation. The difference of adaptation rate of the neurons further assemble features of the input sequences in various timescales. Therefore, given the previous states $S(0), S(1), \dots, S(t)$, their spatiotemporal features will be self-organised on different levels of the network. So the MTRNN is not only a continuous time recurrent neural network that can predict the next states $S(t + 1)$ of the time sequence, but also its internal state acts as a hierarchical memory to preserve the temporal features of the non-linear dynamics in different timescales.

Fig. 2: Language Learning Model based on MTRNN



2.2.1 Learning

In general, the training of the MTRNN follows the updating rule of classical firing rate models, in which the activity of a neuron is determined by the average firing rate of all the connected neurons. Additionally, the neuronal activity is also decaying over time following an updating rule of leaky integrator model. Therefore, when time-step $t > 0$, the current membrane potential status of a neuron is determined both by the previous activation as well as the current synaptic inputs, as shown in Eq. 6:

$$\tau_i u'_{i,t} = -u_{i,t} + \sum_j w_{i,j} x_{j,t} \quad (6)$$

where $u_{i,t}$ is the membrane potential, $x_{j,t}$ is the activity of j -th neuron at t -th time-step, $w_{i,j}$ represents the synaptic weight from the j -th neuron to the i -th neuron and τ is the time scale parameter which determines the decay rate of this neuron. One of the features that is similar to the generic continuous time recurrent neural networks (CTRNN) model is that a parameter τ is used to determine the decay rate of the neural activity; a larger τ means their activities change slowly over time compared with those with a smaller τ .

Assuming the i -th neuron has the number of N connections (i.e. the total number of the neurons in the network is N), Eq. 6 can be transformed into

$$u_{i,t+1} = (1 - \frac{1}{\tau_i})u_{i,t} + \frac{1}{\tau_i} [\sum_{j \in N} w_{i,j} x_{j,t}] \quad (\text{if } t > 0) \quad (7)$$

When the time-step $t = 0$, the membrane potential of the IO neurons is set to 0 and the context neurons are set to initial states:

$$u_{i,0} = \begin{cases} 0, & \text{if } t = 0 \text{ and } i \in I_{IO}, \\ C_{s_c(i,0)}, & \text{if } t \neq 0 \text{ and } i \notin I_{IO} \end{cases} \quad (8)$$

The neural activity of a neuron is calculated in two methods, depending on which level the neuron belongs with:

$$y_{i,t} = \begin{cases} \frac{e^{u_{i,t}}}{\sum_{j \in Z} e^{u_{j,t}}}, & \text{if } i \in I_{IO}, \\ \frac{1}{1 + e^{-u_{i,t}}}, & \text{otherwise.} \end{cases} \quad (9)$$

Thus there is a sigmoid activation function for context neurons, while the input-output neurons are calculated by the soft-max function. The soft-max activation function gives rise to the recovery of a similar probability distribution as the SOM pre-processing modules. Therefore, this activation function results in a faster convergence to the MTRNN network training.

During training, the neurons of MTRNN self-adapt their weight matrices as well as the internal states of the neurons on the context layers for the processing of the incoming time sequence. The purpose of the training is to minimize the error E which is defined by the Kullback-Leibler divergence in this case:

$$E = \sum_t \sum_{i \in O} y_{i,t}^* \log(\frac{y_{i,t}^*}{y_{i,t}}) \quad (10)$$

where $y_{i,t}^*$ is the desired neural activation of the i -th neuron at the t -th time-step, which acts as the target value for the actual output $y_{i,t}$. The target of the training is to minimize E by back-propagation through time (BPTT).

In the BPTT algorithm, the input of the IO neuron is calculated from a mixed partition value r (called the feedback rate) of the previous output value y and the desired value y^* . (Eq. 11)

$$x_{j,t+1} = (1 - r) \times y_{j,t} + r \times y_{j,t}^* \quad (11)$$

where we will use $r = 0.1$ during training, and $r = 0$ during generation, which means that the network is used to generate the sequences autonomously.

$$\frac{\partial E}{\partial u_{i,t}} = \begin{cases} y_{i,t+1} - y_{i,t+1}^* + (1 - \frac{1}{\tau_i}) \frac{\partial E}{\partial u_{i,t+1}}, & \text{if } i \in I_{IO}, \\ \sum_{k \in I_{all}} \frac{\partial E}{\partial u_{i,t+1}} [\lambda_{i,k} (1 - \frac{1}{\tau_i}) + \frac{1}{\tau_k} w_{ki} \int' (u_{i,t})], & \text{otherwise.} \end{cases} \quad (14)$$

In Eq. 12 and Eq. 13, the partial derivatives for w and b are the sums of weight and bias which determine the changes over the whole sequence respectively, and η and β denote the learning rates for the weight and bias changes. Particularly, the term $\partial E / \partial u_{k,t}$ can be calculated recursively as Eq. 14, where the $\int'()$ is the derivative of the Sigmoid Function defined by Eqs. 8 and 9. The term $\lambda_{i,k}$ is the Kronecker’s Delta, whose output is 1 when $i = k$, otherwise it is set to 0.

3 Experiments

To examine the network performance, we recorded the real world training data from object manipulation experiments based on an iCub robot [40]. This is a child sized humanoid robot which was built as a testing platform for theories and models of cognitive science and neuroscience. Mimicking a two-year old infant, this unique robotic platform has 53 degrees of freedom. As such, using the iCub, we set a learning scenario in which a human instructor was teaching the robotic learner a set of language commands whilst providing kinaesthetic demonstration of the named actions. The aim of these experiments was to evaluate the error for generalisation with a large data-set. We were also interested in the mechanisms, especially the neural activities in the hierarchical architecture, which result in such a generalisation.

3.1 Experimental Setup

Fig. 3 shows the setup used in our experiments. The data set was obtained using the following steps:

At the n -th iteration of training, the synaptic weights and the biases of the network of neuron i are updated according to Eq. 12.

$$\begin{aligned} w_{i,j}^{n+1} &= w_{i,j}^n - \eta_{i,j} \frac{\partial E}{\partial w_{i,j}} \\ &= w_{i,j} - \frac{\eta_{i,j}}{\tau_i} \sum_t x_{j,t} \frac{\partial E}{\partial w_{i,t}} \end{aligned} \quad (12)$$

$$b_i^{n+1} = b_i^n - \beta_i \frac{\partial E}{\partial b_i} = b_i - \beta_i \sum_t \frac{\partial E}{\partial u_{i,t}} \quad (13)$$

1. Objects with significantly different colours and shapes were placed at 6 different locations in front of the iCub.
2. A vocal command was spoken by an instructor according to the visual scene that was perceived by the iCub. A complete sentence of the vocal command is composed of a verb and a noun. For instance, assuming we have the command “lift [the] ball”, this was recognised by the speech recognition software called Dragon dictate¹, with which the corresponding verb and noun were recognised and then translated into two dedicated discrete values based on the verb and noun dictionaries (Tab. 1).
3. The built-in vision tracker of the iCub searches for a ball-shaped object based on the dictionary-generated values; the iCub uses its vision tracker system which incorporates visual segmentation algorithm to track a particular type of object, rotate the joints of head and neck and locate it in the visual field.
4. Once the object is located, the iCub rotates its head and triggers the object tracking, which will change the encoder values of the neck and eyes.
5. Joint positions of the head and neck are recorded. The sequence recorder module of the iCub was used to record the sensorimotor trajectories while the instructor was guiding the robot by holding its arms to perform a certain action for each object.
6. The hand and torso joints rotate to certain angles to accomplish the lifting action toward the ball (with human instructor during training/without human instructor during execution)

¹ <http://www.nuance.co.uk/dragon/index.htm>

The whole experimental setup used combinations of 9 actions and 9 objects. From these combinations, both the vocal commands (i.e. a complete sentence includes verb and noun) and the sensorimotor sequences can be created. To the best of our knowledge, this 9×9 noun-verb scenario is one of setups with the highest combination of verbs and nouns in grounded robot language experiments (e.g. [65, 73]) We used such a large number of data to test the combinatorial complexity and mechanical feasibility of this model, as well as to evaluate the generalisation ability and its internal non-linear dynamics when using such a large data-set. From an engineering point of view, after testing the feasibility of generalisation, it is also possible to apply this model in a real-world robot application.

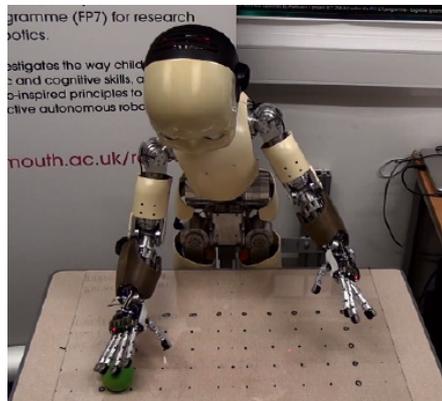
Table 1: Dictionaries of verbs and nouns for the data sets: The instructor showed the robot with different combinations from the 9 actions and 9 nouns. The actions and the objects are represented in two discretised values for semantic command inputs which range from 0 – 0.9. For instance, the command “lift [the] ball” is translated into values [0.8, 0.2].

Actions	Slide Left	Slide Right	Touch	Reach	Push
Verb Value	0.0	0.1	0.2	0.3	0.4
Actions	Pull	Point	Grasp	Lift	
Verb Values	0.5	0.6	0.7	0.8	
Objects	Tractor	Hammer	Ball	Bus	Modi
Noun Value	0.0	0.1	0.2	0.3	0.4
Objects	Car	Cup	Cubes	Spiky	
Noun Values	0.5	0.6	0.7	0.8	

As mentioned before, each speech command was recognised and translated into two semantic command units. Using 9 discretised values for verbs and 9 for nouns, the semantic commands have thus 81 possible combinations. This translation was done according to the verb and noun dictionaries, as shown in Tab. 1. Since we used the visual object tracker in the iCub, the joints of neck and eyes automatically represent the location of the particular object which is presented in the vocal commands. Also the movements of the joint angles in the torso are recorded as the sequences of the motor actions. During the data recording, each recording sequence lasted 5 second and the encoder values of 41 joints were sampled at 50ms intervals. Thus, the complete input vector of the data set contains 100 temporal steps of the discrete semantic command, location of visual attention and joint movement of the torso, as shown in Tab. 2.

Three experiments were carried out and are described in the next subsections: in the first experiment, given the 9 actions and 9 objects data set, we will search the parameter space and find the best parameters for the network training.

Fig. 3: Experimental Scenario



In the second experiment, the training and generalisation performance will be shown given different types of manipulated data sets. For the third experiment, we will further analyse the generalisation ability of the MTRNN network. All these experiments were run using a modified version of the Aquila software [51] in a GPU computer with one Tesla C2050 and two GeForce GTX 580 graphic cards.

3.2 Experiment 1 - Training Setting

During this section, we used the data set consisting of the complete 9×9 combinations (i.e. $N_v = 9, N_n = 9$), which include information about 6 different object locations. Thus the whole data-set contains $9 \times 9 \times 6 = 486$ sequences, which were all used for training the network. The most distinct feature of the MTRNN, with respect to generic RNN or CTRNN networks, is that different neurons have distinct time constants τ , which are also one of the key factors that determine the training performances of the network. In this experiment, we systematically changed these parameters in the parameter space in order to find out the best parameter settings.

A parameter space is defined as $(\tau_s, \tau_f, N_{C_s}, N_{C_f})$, representing the time constants on the context slow layer, context fast layer and the number of neurons on these two layers respectively. In order to minimise the effects of the randomness of the initialisation, a total of 3 training trials were done with the same training setting, as shown in Tab.3. Previous MTRNN experiments in the literature have reported different parameters settings ([73], [24] and [27]). From these experiments, we can discover that the number of neurons on the C_s and C_f layers were determined mainly according to dimension of the IO layer, but they generally kept a ratio from 1 : 4 to 1 : 3. To start, we firstly set the parameters according to the minimum values (70, 5, 20, 60) from previous research [73]. Then we scaled up the numbers of neurons

Table 2: Structure of the Training Data

Description	Semantic Commands	Object Location (Neck and Eyes)	Torso Joints
Dimension	2	6	3
Description	Left Arm Joints	Right Arm Joints	
Dimension	16	16	

Table 3: Training Error with Different Parameter Settings ($C_s, C_f, N_{C_s}, N_{C_f}$)

Parameters	Error 1	Error 2	Error 3	Ave.
(70, 5, 20, 60)	0.084	0.081	0.085	0.0833
(70, 3, 20, 60)	0.084	0.085	0.082	0.0837
(70, 5, 30, 60)	0.084	0.086	0.083	0.0843
(70, 5, 30, 50)	0.082	0.079	0.080	0.0803
(70, 5, 30, 100)	0.079	0.079	0.078	0.0787
(70, 5, 60, 100)	0.078	0.078	0.077	0.0777
(70, 5, 40, 120)	0.079	0.079	0.078	0.0787
(70, 5, 50, 140)	0.075	0.075	0.077	0.0757
(70, 5, 60, 160)	0.072	0.071	0.074	0.0723
(70, 5, 50, 120)	0.071	0.070	0.071	0.0707
(70, 3, 50, 120)	0.071	0.071	0.070	0.0707
(70, 5, 70, 120)	0.070	0.071	0.072	0.0710

on the context layers and adjusted the time constants. The less crucial parameters were kept constantly: learning rates $\eta = 0.7$, $\beta = 0.7$, momentum = 0.9, weightRange = 0.025. The stopping criteria for the training process was that the error did not decrease more than 1^{-6} within consecutive 100 iterations. From Tab. 3, we can see that the number of neurons on the context layers affected much on the training performance of the network. Comparatively, the time constants played a less significant role for the training error than the number of neuron did. Also results showed that the suitable ratio for numbers of C_s and C_f neurons should be kept to around 1 : 4 to 1 : 3.

3.3 Experiment 2 - Training Performance

From the previous experiment, we found that the best parameters for training the 9×9 verb-noun data-set are those shown in bold in Tab. 3, among which we selected (50, 5, 70, 100). We then examined the training performance of the network under this parameter setting using different data-sets. To test the generalisation ability, these data-sets were manipulated: a subset of the combinations of actions and objects were removed from the training set, to be used as validation test sets when testing the generalisation ability of the network. The detailed information about the manipulated data-sets are shown in Tab. 4, where the coloured numbers N indicate the specific verb-noun combination removed in the specific N -th data-set. We can see that the number of removal sets were increasing from the first to the third test-set, indicating the difficulty of generalisation was increasing. Also at the second and the third data-sets, some of the removal sets were

next to each other, which further increased the difficulty of generalisation.

Using the parameter set of (50, 5, 70, 100), the training curves (Figs. 4) show that the training converged. To further demonstrate the robustness of the generalisation ability given the untrained sensorimotor sequences, the validation sets, which were not included in the training, were fed into the network. In this way we aimed to test how the network responds to noun-verb combinations not used during training. Using the three MTRNNs we trained from three data-sets, we performed three generalisation experiments using the missing verb-noun combinations. In the experiments, only the first time step data in the sequence was provided (i.e. $r = 0$ in Eq. 11), which includes the initial position of the torso, head and eye motors, as well as the vocal command. Then the network prediction was used as the input of the next time-step and formed a closed-loop to complete 100-step of the time sequence generation. The errors of the whole three training-sets, as well as those in different steps are shown in Tab. 5. A more straightforward visualisation of the network performance can be found in Figs. 5, which displays three examples of generated time sequences for motor actions from three MTRNNs. As we calculated in Tab. 5, the training error became larger when the number of training samples was smaller. In particular, a larger error could be found at the beginning of each time sequence, but the network became stable and generated a stable motor trajectory with less error as time elapsed. There were some errors displayed in the trajectories generation, so sometimes the generated robot behaviours based on the trajectories are biased with the orig-

Table 4: Some of the sequences containing particular semantic combinations of verbs and nouns were removed during training. The number i in the cell indicates that such a combination was removed in the i -th training set for generalisation experiments.

V. \ N.	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8
0.0	1/2/3			3		2	3		
0.1	2/3	1		3		2	3		
0.2	3	2	1	3			2/3		
0.3		2/3		1	3		2	3	
0.4		3	2		1/3			2/3	
0.5		3	2		3	1		2/3	
0.6			3	2		3	1		2/3
0.7			3	2		3		1	2/3
0.8			3		2	3			1/2/3

Table 5: RMS Error of the Generalisation Tests

Test	1	2	3
RMS Error (All)	0.0052	0.0069	0.0169
RMS Error (Step 1-20)	0.0064	0.0082	0.0240
RMS Error (Step 21-40)	0.0042	0.0075	0.0194
RMS Error (Step 41-60)	0.0033	0.0069	0.0150
RMS Error (Step 61-80)	0.0031	0.0062	0.0121
RMS Error (Step 81-100)	0.0024	0.0052	0.0101

inal ones. However, in most of cases, the generated robot behaviours correctly followed the semantic commands ².

4 Generalisation Analyses

In this section, we focus on the problem of how the verb-noun generalisation ability of the MTRNN network is achieved. As shown in the previous section, the network was able to “understand” the un-trained verb-noun combinations in the sense that the generated time sequences for motor actions were close to the originals, so the iCub robot can perform the action named in the vocal commands. For an experiment with a similar aim, [62] reported combining two hierarchical recurrent neural networks which can also accomplish verb-noun generalisation for understanding combinatorial semantics in a situated environment. The model they used, called recurrent neural networks with parametric biases units (RN-NPB), had similar non-linear dynamics as the MTRNN: the non-linear dynamics are determined by a small number of neural units which act as bifurcation for the whole system. Particularly, in our case, the learning sequences contain a much larger dimension (35) of the motor joint angles for the iCub movements, compared with motor sequences that trained in [62]. These complex sequences result in the bifurcation which occurs hierarchically in the MTRNN structure.

From this point, we hypothesise that the MTRNN, or any other hierarchical RNNs, results in the separation in the net-

work dynamics about different modalities in a self-organised way along the lexicon categories of vocal commands after training. The type of separation depends on the different organisation of the training data structures, This separation, with the constraints of sensorimotor sequences, occurs on different levels of the hierarchical architecture using different strategies. The way the network presents such separation in the hierarchical dynamics is self-organised, and largely depends on the data structure of the training data in the spatial and temporal domains. Particularly in our experiment setting, after enough training, the synaptic weights between a *basic* motor behaviour are enforced with a particular dimension about the verb input, as it dominates a large portion of the spatio-temporal space in the sensorimotor sequences. The *basic* motor behaviour here means that such kind of motor actions belong to general definitions such as “slide”, “touch” and etc, without a specific goal for directing action. This is similar to the mechanism that the hearing of a verb causes in our brain: a specific area in the pre-motor cortex, corresponding to certain motor action fires when a particular verb is heard or said. On the contrary, the noun also affects part of the sensorimotor outputs in terms of its role to offset the basic motor behaviours into a specific goal-directed action.

In the following experiments, we will examine this hypothesis by means of manipulating data and visualising the training results.

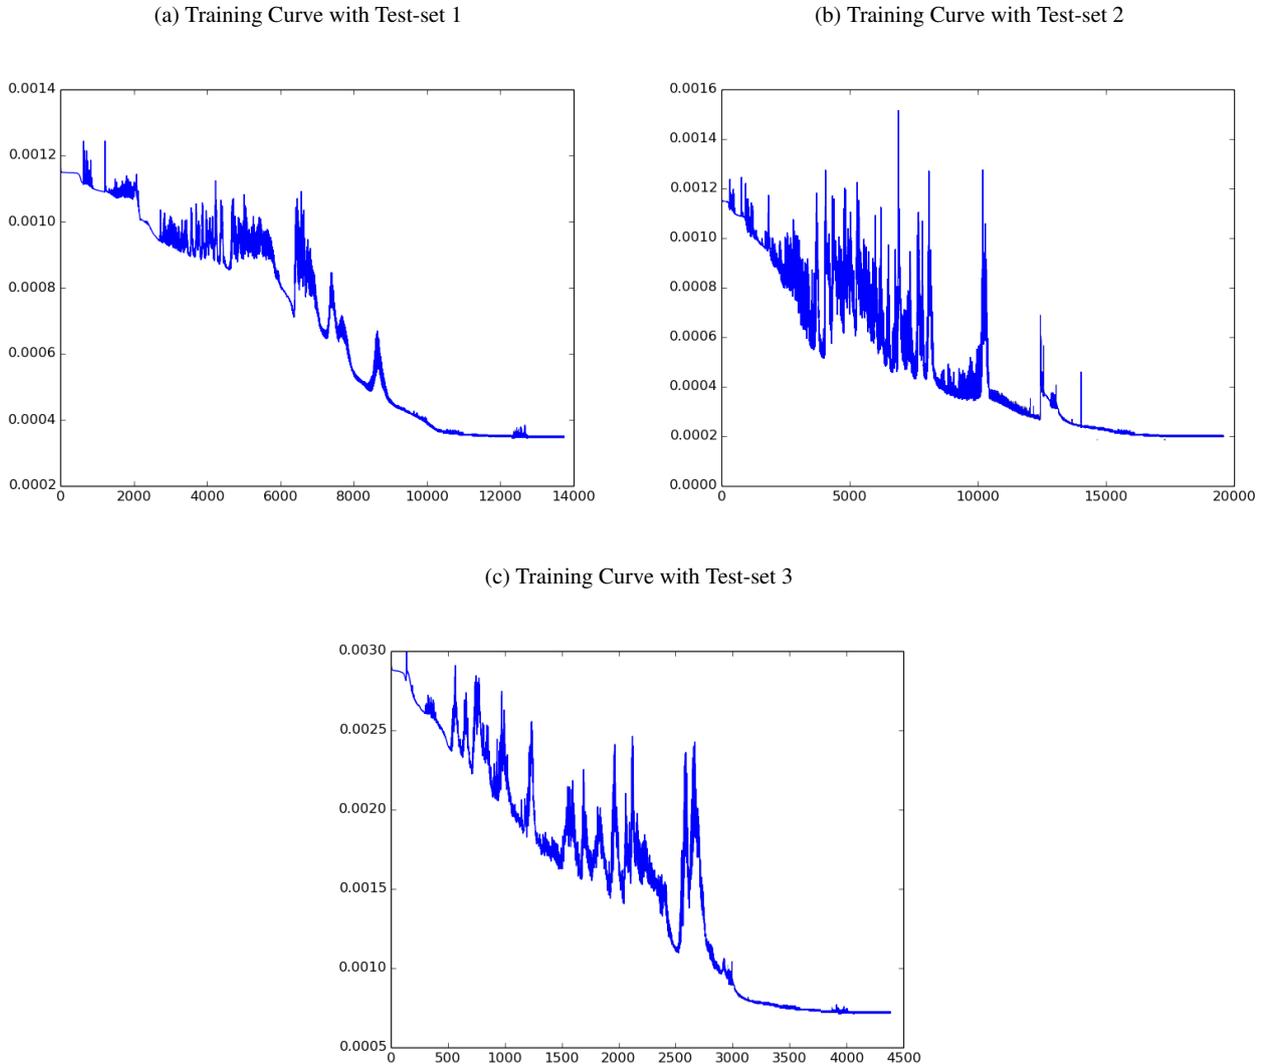
4.1 Generalisation with Partial Inputs

In this subsection, we concentrate on the comparisons of the results after the removal of different modalities. These comparisons included two parts: i) Error of generalisation after removals; ii) Visualisation of weights after removals.

For the first part of the analysis, in order to obtain a more conclusive statement, we used two sets of data 9×9 and 3×3 of verb-noun combinations. The 3×3 data-set (Tab. 7) contains all the combinations of three actions and three objects, which were placed into 6 different locations. Tab. 6 was

² <https://youtu.be/FOgKbJ-iEhM>

Fig. 4: Training Curves with Three Test-sets



used for the vocal command discretisation. For the second part of the experiment, the visualisation of weights was only done with the 3×3 data-sets, since its features are easier to observe and its basic principle can be easily extended to the 9×9 data-set.

Table 6: Dictionary of verbs and nouns for the 3×3 data sets

Actions	Slide Left	Slide Right	Touch
Objects	Tractor	Hammer	Spikey
Values	0.1	0.2	0.3

For both parts of the experiment, in order to observe how different lexical categories and visual input affected the training

results, especially within the output of the sequences of the motor behaviours, different parts of the input data were removed:

1. No modification (base-line)
2. Remove the noun input (i.e. the first input unit was reset to zero.)
3. Remove the verb input (i.e. the second input unit was reset to zero.)
4. Remove the location of the visual object (i.e. from the third to eighth units were reset to zero.)

During the generalisation tests, the full un-trained data was placed into the network. The training and generalisation error of the motor output was compared in the Tab. 8 and Tab. 9. From these two tables we can see that the removal of the verb resulted in a larger generalisation error than the other

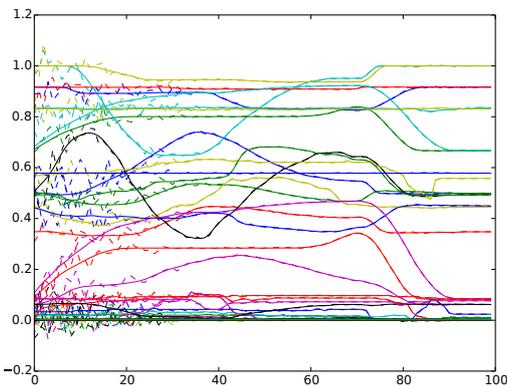
Fig. 5: Trajectory Generation

The generated trajectories (dotted) with 41 dimensions were plotted and compared with the original trajectories. Three test-sets were selected to validate the training performances with different training sets.

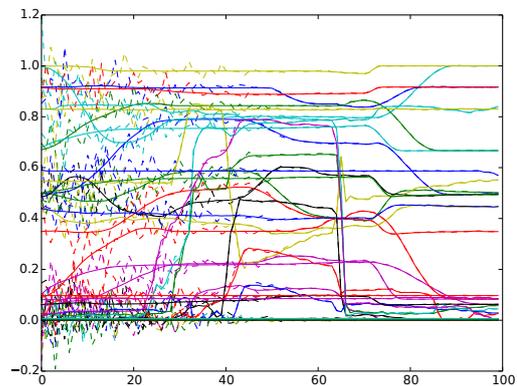
(a) Generated Trajectory from MTRNN 1, Test-set 61 (v.-n.: 0.1-0.1)



(b) Generated Trajectory from MTRNN 2, Test-set 231 (v.-n.: 0.4-0.2)



(c) Generated Trajectory from MTRNN 3, Test-set 484 (v.-n.: 0.8-0.8)



two tests, while the removal of the object location resulted in the lowest generalisation error.

For the second part of the experiment, the main aim was to understand the effect of a particular input modality (e.g. lexical structure and visual input) in the whole network training, by observing the visualization of the weights. We conducted an experiment with a smaller data-set (3×3) than the previous experiments, due to the fact that smaller number of weights give a better presentation for the visualization. But a similar conclusion would be extended into the larger 9×9 data-set. Figs. 6 visualised the weighting matrix, where the neurons from number 0 to number 703 were neurons on

the IO layer, from number 704 to number 764 were neurons on the C_f layer and from number 765 to number 794 were neurons on the C_s layer. The weight matrices in Fig. 6a, Fig. 6c and Fig. 6d looked quite similar. But in Fig. 6b, without the verb input, we could easily notice that a large amount of weights from IO layer to C_f remain to be untrained. To quantitatively evaluate this observation, Tab. 10 calculated the 2-norm to obtain the Euclidean distances from the manipulated weighting matrices to the base-line matrix. The

Table 7: Removal of data in the 3×3 data-set. The number i in the cell indicates that such a combination was removed in the i -th training set for the generalisation experiments.

	0	1	2
0	1		2
1		1/2	
2	2		1

Table 8: Removal Part of Input (3 verbs and 3 nouns)

Error	w/o verb training	w/o verb generalisation	w/o visual location training	w/o visual location training	w/o noun training	w/o noun generalisation
Test 1	0.0003	0.1041	0.0003	0.0594	0.0003	0.0868
Test 2	0.0003	0.1129	0.0003	0.0612	0.0003	0.0933

Table 9: Removal Part of Input (9 verbs and 9 nouns)

Error	w/o verb training	w/o verb generalisation	w/o noun training	w/o noun generalisation	w/o visual location training	w/o visual location training
Test 1	0.0003	0.5311	0.0003	0.5223	0.0003	0.0921
Test 2	0.0005	0.6623	0.0005	0.7473	0.0005	0.1379
Test 3	0.0006	0.8574	0.0006	0.7494	0.0006	0.1771

2-norm was calculated by:

$$d(\mathbf{W}^m - \mathbf{W}^b) = \sqrt{\sum_{i=1}^n \sum_{j=1}^n (d_{ij}^m - d_{ij}^b)^2} \tag{15}$$

where \mathbf{W}^m is the weighting matrix after data manipulation, \mathbf{W}^b is the weighting matrix from base-line experiment, d is the weight from the i -th neuron to j -th neuron. Here $n = 795$ which is the total number of neurons.

From the comparisons of weight matrices and the Euclidean distances, we further verified our hypothesis that the lexical structure of verbs plays a significant role in the training, since it is further grounded in the differences of motor action trajectories, which dominate a large spatio-temporal space of the sequences.

Table 10: Euclidean Distances between Partial Input Matrices and Normal Training Matrix

	w/o Verb	w/o Noun	w/o Location
Distance	8.9100	0.9450	0.6736

4.2 Internal Dynamics

In the previous analysis, we have looked at the generalisation ability of the MTRNN. A preliminary conclusion suggests

that the lexical structure of the verb plays a significant role in maintaining the convergence of the temporal sensorimotor sequences. In this section we are particularly interested in how the generalisation capabilities are brought by the recurrent connected hierarchical structure. We believed that part of these answers can be found by observing the detailed neural activities on each context layer given the selection of different inputs. The neural activities were therefore examined using the 9×9 data-set, with a previously trained MTRNN with the parameter setting of (50, 5, 70, 100).

The following figures showed the PCA trajectories of the internal neural dynamics on the C_f (Fig. 7, Fig. 8 and Fig. 9) and C_s (Fig. 10, Fig. 11 and Fig. 12) layers. Since the complete 9×9 data-set contains 486 sequences, whose patterns can hardly be observed in one single figure, only a few samples were presented in the following figures to clearly show the PCA trajectories. Fig. 7 and Fig. 10 showed the selected PCA trajectories on the C_f and C_s layers. These trajectories mainly concern combinations of verb inputs and a few noun inputs. We can see that the verbs mainly determine the patterns of the trajectories, which implies that the processing of verbs mainly affects the temporal dynamics in the MTRNN.

The following figures mainly show how the differences in lexical structures and visual information result in the differences in the PCA trajectories. Fig. 8 and Fig. 11 show the PCA trajectories of the internal dynamics on C_f and C_s layers, with different noun inputs; Fig. 9 and Fig. 12 showed the PCA trajectories with different object location inputs. We could observe that the differences of nouns on the C_f (Fig. 8)

cause divergences at the beginning of the trajectories, but not at the end. From Fig. 9 comparisons show the differences of visual inputs produce even smaller divergences in the trajectories, and that the divergences mainly occurred at the middle of the trajectories. Comparatively, from the activities on the C_s layer (Fig. 11 and Fig. 12), the divergences of the trajectories from nouns and visual inputs were even smaller: the C_s layer mainly encoded the information from the verbs.

To summarise the MTRNN analysis, the model self-organises similar patterns on various levels for every sensorimotor sequence, reflecting the hierarchical structure for the vocal commands. Particularly, we can see that the difference between verb inputs results in larger divergence of the trajectories than noun and object-location differences. Due to the data structure of our input vectors, the IO layer represents a collection of each word. With a slower adaptation rate than the IO layer, the C_f represents the grounded meaning of each verb, noun and visual information. This grounding process is learnt by all temporal sensorimotor sequences. Similarly, using slower changing neurons, the C_s layer represents the general motor behaviour (i.e. the verb) of the whole sensorimotor sequence.

Therefore, the C_f activation mainly represents the lexical structures (verbs and nouns). The visual location has a limited effect on the C_f activation, probably because the information of noun already has overlap with the object information about visual location. As the main factor of the C_f layer, the same verbs are represented as a similar pattern on the fast context layer in all Fig. 7, Fig. 8 and Fig. 9. The difference from nouns can be observed at the beginning of the trajectories. It may correspond to the difference of robot behaviours at the beginning of the time sequences, caused by the neck and eye tracking before the actual hand movement starts. Comparing with the C_f layer, the C_s activation changes even slower. It generally represents the motor behaviours; only the verbs are represented in different patterns.

5 Discussion

5.1 Functional Hierarchy of RNN and its Bifurcation

It has been reported that quite a few RNN models based on functional hierarchy, such as RNNPB, MTRNN and conceptors [30], allow the bifurcation to occur in the RNN dynamics. We will give a brief discussion of how this bifurcation happens. Assuming we have a simple hierarchical RNN with an additional unit (which can be regarded as a simplified version of RNNPB) as depicted in Fig. 13. The system can be described as Eq. 16.

$$\begin{cases} \dot{x}_1(t) = & -x_1(t) + f(x_3(t)) \\ \dot{x}_2(t) = & -x_2(t) + a \cdot f(x_1(t)) + c \cdot PB \\ \dot{x}_3(t) = & -x_3(t) + b \cdot f(x_2(t)) \\ y(t) = & f(x_3(t)) \end{cases} \quad (16)$$

There are three fixed points in this network. After the network has been trained, i.e. the weights a , b and c are fixed, the coordinates of fixed points only depends upon the value of PB . Furthermore, the coordinates of the fixed points $[x_1, x_2, x_3]$ are first-order functions of the value of PB units (please see appendix for the calculation in details). In other words, the coordinates of the fixed points further determine the domain of different bifurcation properties. This is the reason that changing the parameter of PB units will change the qualitative structure of the non-linear dynamics of the network. From the bifurcation explanation of the simplified RNNPB model, at the next step we can also extend this to other hierarchical RNNs such as MTRNN, as they are holding a fundamentally similar theoretical foundation [64].

5.2 Generalisation Ability of MTRNN

In our experiments, the MTRNN was trained under a particular input data structure: Firstly the language commands were recorded as auditory data and transformed into a discrete symbolic representation, and secondly the object locations and the motor behaviours were also stored as the angles of motor joints. This unique structure is a simplified representation of the common coding theory, which proposes that perceptual inputs and motor actions are sharing the same format of the representation within the cognitive processes.

The neural dynamics in our MTRNN exhibited a dynamics which are different from those reported in [27] and [23]. Whereas the noun (or object perceptual inputs) play a significant factor in the dynamics of context layers in these two examples, our network has minimised the effects of nouns or the object perception. This is partly because the input data structure, where the motor joints of the iCub robot have much larger dimensions than the visual perception input. Also the spatial information for objects in our experiment setting is much easier to learn, compared with our diversified motor behaviours. The generalisation here concerns more the inference of symbolic meaning of a language command due to the composition of neural dynamics. During the training in a hierarchical network, such as MTRNN or RNNPB, the neural connections strengthen between a particular type of sensorimotor sequence (motor angle changes due to differing behaviours) and visual perception. Particularly, in our case of 9×9 data-sets, most of our network weights stores the learning of motor actions.

Note that the generalisation of commands in the verb-noun combinations is not the same as we usually do in the generic recurrent neural networks (e.g. [29,52,77]), which expect the network to do interpolation or extrapolation with a novel input value in either temporal or spatial space. While generalizing dynamical patterns by interpolation is a non-trivial task for training motor patterns in robots, our main concern is the novel combinations in the context of lexicon acquisition. In our case, the learning of verbs and nouns results in the emergence of different dynamics that are mostly stored in different synaptic weights, and thus their combinatorial composition is realised by the non-linearity of the recurrent connections. Considering the different generalisation abilities of generic RNN, RNNPB [33,77] and MTRNN [23], the hierarchical RNNs appear particularly suitable for the production of flexible motor behaviour and language expression simultaneously in the real-world social robot experiments.

$$\prod_{t=1}^{T'} P(y_t | y_{t-1}, y_{t-2}, \dots, y_1, c) = P(y_{T'}, y_{T'-1}, \dots, y_1 | x_T, x_{T-1}, \dots, x_1) \tag{17}$$

Generic RNNs are not able to approximate the probability of the sequence with arbitrary length because of its vanish gradient problem, but other novel RNNs, such as LSTM, BRNN (Bi-directional Recurrent Neural Networks), have been successfully employed to construct the ED architecture to “understand” (encode) and to “generate” (decode) the temporal sequences. Furthermore, due to the recent popularity of parallel computation by GPU, it has become possible to train and use such architectures to solve problems such as machine translation and image captioning.

As the MTRNN can also avoid the vanish gradient problem, and larger MTRNN can be implemented via GPU, it is also possible to embed the MTRNN into the ED architecture. In fact, the context slow level C_s already exhibits a similar feature of “thought vectors”, using a stable neural vector to represent the basic profiles of motor actions and object instances (in our robotic experiment). They also have similar information bi-directional flows which allow the networks to recognise and to generate the time sequences. Despite their similarities, compared with LSTM, the MTRNN have other distinct features: First, from the above experiments and from other MTRNN experiments [23,27], it has been shown that the fast context layers and slow context layers exhibit various dynamics to explicitly represent the relationship between the verbs and nouns. The deep LSTM, on the contrary, has not been reported to have similar dynamics. Second, differ-

5.3 Thought Vectors and Further Development

A few machine learning methods have recently been proposed based on the encoder-decoder (ED) architecture [12], which achieved great performance in machine translation [63], image captioning [69], etc. The ED architecture usually consists of two recurrent neural networks. One deep RNN network encodes a sequence of input vectors with arbitrary length into a fix-length vector representation in a hierarchical way, while the other deep RNN network decodes this representation into a target sequence of output vector. This specific representation between the encoder and the decoder RNNs is called “thought vectors” which is claimed to represent the meaning of the sequence in a high-dimensional space. The training of such an architecture is done by maximizing the conditional probability of the target sequence. If the input sequence is denoted as (x_1, x_2, \dots, x_T) and the corresponding output sequence is $(y_1, y_2, \dots, y_{T'})$ (T does not necessarily equal to T'), the next symbol generation is done by maximising Eq. 17.

ently from the static vector representation from LSTM, the context layers allow a “slow” change through time which is more realistic for an interaction environment, where it can be used to dynamically exhibit the meaning of sentences and sensorimotor information.

Admittedly, the training of deep RNNs, e.g. LSTMs and MTRNNs, costs a large amount of computational effort. But the recent development of GPU computing provides an opportunity to construct and test such a big scale neural network with a reasonable time and budget. The combination of MTRNN, the concept of “thought vectors” and its embodiment in robotic systems, will allow us to further explore issues such as:

1. The comparison of the performances of MTRNN, LSTM and BRNN within the ED architecture and examine their performances in the robotic platforms.
2. The robot motor action, as a natural temporal sequence, can be further incorporated as the training of RNNs of ED architecture with connections to other modalities.

6 Conclusion

This paper presents a neurobotic study on noun and verb generation and generalisation, utilising with the MTRNN networks, with a large data-set, consisting of vocal language commands, visual object and motor action data. Although

the generalisation abilities of hierarchical RNNs (RNNPB, MTRNN) have been reported in previous research, this is the first study to demonstrate its generalisation capability using such a large data-set, which enables the robot to learn to handle real-world objects and actions. These experiments showed that the generalisation ability of the network are possible even with a large amount of test-sets (9 motor actions and 9 objects placed in 6 different locations). This is particularly important because the recurrent connections between the verbs and nouns are associated with different modalities of the training-data, which is strengthened during embodiment training by the sensorimotor interaction. Detailed analyses on the robot’s neural controller showed that the dynamics on different layers are self-organized in the MTRNN. These self-organised dynamics further constitute a functional hierarchical representation on different layers, which associate different lexical structures with different modalities of the sensorimotor inputs. The MTRNN showed how the embodied information about the verbs dominates a large portion of the network dynamics, since the proprioception information plays a significant role in the training sequences. As such, the hierarchical RNNs, such as MTRNN, are shown to be particularly beneficial in building a neuro-robotics cognitive architecture about language learning for robotic systems, where the recurrent connections are able to self-organise and build associations between embodied information in different modalities and the lexical structure information.

Acknowledgements This research has been supported by the EU project POETICON++ under grant agreement 288382 and the UK EPSRC project BABEL. We are grateful to Dr. Christopher Ford for his helpful review.

References

1. Minoru Asada, Koh Hosoda, Yasuo Kuniyoshi, Hiroshi Ishiguro, Toshio Inui, Yuichiro Yoshikawa, Masaki Ogino, and Chisato Yoshida. Cognitive developmental robotics: a survey. *Autonomous Mental Development, IEEE Transactions on*, 1(1):12–34, 2009.
2. Hiromitsu Awano, Shun Nishide, Hiroaki Arie, Jun Tani, Toru Takahashi, Hiroshi G Okuno, and Tetsuya Ogata. Use of a sparse structure to improve learning performance of recurrent neural networks. In *Neural Information Processing*, pages 323–331. Springer, 2011.
3. Elizabeth Bates and Frederic Dick. Language, gesture, and the developing brain. *Developmental psychobiology*, 40(3):293–310, 2002.
4. Cynthia L Breazeal. *Designing sociable robots*. MIT press, 2004.
5. Angelo Cangelosi. Grounding language in action and perception: from cognitive agents to humanoid robots. *Physics of life reviews*, 7(2):139–151, 2010.
6. Angelo Cangelosi, Giorgio Metta, Gerhard Sagerer, Stefano Nolfi, Christopher Nehaniv, Kerstin Fischer, Jun Tani, Tony Belpaeme, Giulio Sandini, Francesco Nori, et al. Integration of action and language knowledge: A roadmap for developmental robotics. *Autonomous Mental Development, IEEE Transactions on*, 2(3):167–195, 2010.
7. Angelo Cangelosi and Domenico Parisi. How nouns and verbs differentially affect the behavior of artificial organisms. In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society, London: LEA*, pages 170–175, 2001.
8. Angelo Cangelosi and Domenico Parisi. The processing of verbs and nouns in neural networks: Insights from synthetic brain imaging. *Brain and Language*, 89(2):401–408, 2004.
9. Angelo Cangelosi and Domenico Parisi. The processing of verbs and nouns in neural networks: Insights from synthetic brain imaging. *Brain and Language*, 89(2):401–408, 2004.
10. Angelo Cangelosi and Matthew Schlesinger. *Developmental robotics: From babies to robots*. MIT Press, 2015.
11. Nancy Chang, JEROME Feldman, and Srinu Narayanan. Structured connectionist models of language, cognition and action. *Progress in Neural Processing*, 16:57, 2005.
12. Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
13. Kerstin Dautenhahn. Socially intelligent robots: dimensions of human–robot interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480):679–704, 2007.
14. Kerstin Dautenhahn and Angelo Cangelosi. Progress on evolution of communication and interaction studies. *Interaction Studies*, 2012.
15. Peter F Dominey. Recurrent temporal networks and language acquisition—“from corticostriatal neurophysiology to reservoir computing”. *Frontiers in psychology*, 4, 2013.
16. Peter Ford Dominey, Anthony Mallet, and Eiichi Yoshida. Real-time spoken-language programming for cooperative interaction with a humanoid apprentice. *International Journal of Humanoid Robotics*, 6(02):147–171, 2009.
17. Kathleen M Eberhard, Michael J Spivey-Knowlton, Julie C Sedivy, and Michael K Tanenhaus. Eye movements as a window into real-time spoken language comprehension in natural contexts. *Journal of psycholinguistic research*, 24(6):409–436, 1995.
18. Maxwell Forbes, Rajesh PN Rao, Luke Zettlemoyer, and Maya Cakmak. Robot programming by demonstration with situated spatial language understanding. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 2014–2020. IEEE, 2015.
19. Dedre Gentner. Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. *Center for the Study of Reading Technical Report; no. 257*, 1982.
20. Bruno Golosio, Angelo Cangelosi, Olesya Gamotina, and Giovanni Luca Masala. A cognitive neural architecture able to learn and communicate through natural language. *PLoS one*, 10(11):e0140866, 2015.
21. Frank H Guenther. Cortical interactions underlying the production of speech sounds. *Journal of communication disorders*, 39(5):350–365, 2006.
22. D Geoffrey Hall and Sandra R Waxman. Assumptions about word meaning: Individuation and basic-level kinds. *Child Development*, 64(5):1550–1570, 1993.
23. Stefan Heinrich, Sven Magg, and Stefan Wermter. Analysing the multiple timescale recurrent neural network for embodied language understanding. In *Artificial Neural Networks*, pages 149–174. Springer, 2015.
24. Stefan Heinrich, Cornelius Weber, and Stefan Wermter. Adaptive learning of linguistic hierarchy in a multiple timescale recurrent neural network. In *Artificial Neural Networks and Machine Learning—ICANN 2012*, pages 555–562. Springer, 2012.
25. Xavier Hinaut and Peter Ford Dominey. Real-time parallel processing of grammatical structure in the fronto-striatal system: a recurrent network simulation study using reservoir computing. *PLoS one*, 8(2):e52946, 2013.

26. Wataru Hinoshita, Hiroaki Arie, Jun Tani, Hiroshi G Okuno, and Tetsuya Ogata. Emergence of hierarchical structure mirroring linguistic composition in a recurrent neural network. *Neural Networks*, 24(4):311–320, 2011.
27. Wataru Hinoshita, Tetsuya Ogata, Hideki Kozima, Hisashi Kanda, Toru Takahashi, and Hiroshi G Okuno. Emergence of evolutionary interaction with voice and motion between two robots using rnn. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 4186–4192. IEEE, 2009.
28. Ian S Howard and Piers Messum. Modeling the development of pronunciation in infant speech acquisition. *Motor Control*, 15(1):85–117, 2011.
29. Masato Ito and Jun Tani. Generalization in learning multiple temporal patterns using rnnpb. In *Neural Information Processing*, pages 592–598. Springer, 2004.
30. Herbert Jaeger. Controlling recurrent neural networks by conceptors. *arXiv preprint arXiv:1403.3369*, 2014.
31. Michael P Kaschak, Carol J Madden, David J Theriault, Richard H Yaxley, Mark Aveyard, Adrienne A Blanchard, and Rolf A Zwaan. Perception of motion affects language processing. *Cognition*, 94(3):B79–B89, 2005.
32. Alan W Kersten. An examination of the distinction between nouns and verbs: Associations with two different kinds of motion. *Memory & cognition*, 26(6):1214–1232, 1998.
33. Jens Kleesiek, Stephanie Badde, Stefan Wermter, and Andreas K Engel. Action-driven perception for a humanoid. In *Agents and Artificial Intelligence*, pages 83–99. Springer, 2013.
34. Teuvo Kohonen. The self-organizing map. *Neurocomputing*, 21(1):1–6, 1998.
35. David Landy, Colin Allen, and Carlos Zednik. A perceptual account of symbolic reasoning. *Frontiers in psychology*, 5, 2014.
36. Emiddia Longobardi, Pietro Spataro, Diane L. Putnick, and Marc H Bornstein. Noun and verb production in maternal and child language: Continuity, stability, and prediction across the second year of life. *Language Learning and Development*, pages 1–16, 2015.
37. Mandy J Maguire, Kathy Hirsh-Pasek, and Roberta Michnick Golinkoff. A unified theory of word learning: Putting verb acquisition in context. *Action meets word: How children learn verbs*, page 364, 2006.
38. Olivier Mangin and Pierre-Yves Oudeyer. Learning to recognize parallel combinations of human motion primitives with linguistic descriptions using non-negative matrix factorization. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 3268–3275. IEEE, 2012.
39. Cynthia Matuszek, Evan Herbst, Luke Zettlemoyer, and Dieter Fox. Learning to parse natural language commands to a robot control system. In *Experimental Robotics*, pages 403–415. Springer, 2013.
40. Giorgio Metta, Giulio Sandini, David Vernon, Lorenzo Natale, and Francesco Nori. The icub humanoid robot: an open platform for research in embodied cognition. In *Proceedings of the 8th workshop on performance metrics for intelligent systems*, pages 50–56. ACM, 2008.
41. Marco Mirolli and Domenico Parisi. Language as a cognitive tool. *Minds and Machines*, 19(4):517–528, 2009.
42. Marco Mirolli and Domenico Parisi. Towards a vygotskian cognitive robotics: the role of language as a cognitive tool. *New Ideas in Psychology*, 29(3):298–311, 2011.
43. Dipendra K Misra, Jaeyong Sung, Kevin Lee, and Ashutosh Saxena. Tell me dave: Context-sensitive grounding of natural language to manipulation instructions. *Proceedings of Robotics: Science and Systems (RSS), Berkeley, USA*, 2014.
44. Alva Noë. Experience and the active mind. *Synthese*, 129(1):41–60, 2001.
45. Rony Novianto. Flexible attention-based cognitive architecture for robots. 2014.
46. Takaaki Ogata and Hiroshi G Okuno. Integration of behaviors and languages with a hierarchical structure self-organized in a neuro-dynamical model. In *Robotic Intelligence In Informationally Structured Space (RiSS), 2013 IEEE Workshop on*, pages 89–95. IEEE, 2013.
47. Bruno A Olshausen and David J Field. Sparse coding with an over-complete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325, 1997.
48. Pierre-Yves Oudeyer and Frédéric Kaplan. Discovering communication. *Connection Science*, 18(2):189–206, 2006.
49. Domenico Parisi and Angelo Cangelosi. A unified simulation scenario for language development, evolution and historical change. In *Simulating the evolution of language*, pages 255–275. Springer, 2002.
50. Diane Pecher, René Zeelenberg, and Lawrence W Barsalou. Verifying different-modality properties for concepts produces switching costs. *Psychological Science*, 14(2):119–124, 2003.
51. Martin Peniak, Anthony Morse, Christopher Larcombe, Salomon Ramirez-Contla, and Angelo Cangelosi. Aquila: An open-source gpu-accelerated toolkit for cognitive and neuro-robotics research. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 1753–1760. IEEE, 2011.
52. Fernando J Pineda. Generalization of back-propagation to recurrent neural networks. *Physical review letters*, 59(19):2229, 1987.
53. Friedemann Pulvermüller. *The neuroscience of language: on brain circuits of words and serial order*. Cambridge University Press, 2002.
54. Richard A Reale and Thomas J Imig. Tonotopic organization in auditory cortex of the cat. *Journal of Comparative Neurology*, 192(2):265–291, 1980.
55. Katharina J Rohlfing, Jannik Fritsch, Britta Wrede, and Tanja Jungmann. How can multimodal cues from child-directed interaction reduce learning complexity in robots? *Advanced Robotics*, 20(10):1183–1199, 2006.
56. Ayse Pinar Saygin, Stephen McCullough, Morana Alac, and Karen Emmorey. Modulation of bold response in motion-sensitive lateral temporal cortex by real and fictive motion sentences. *Journal of cognitive neuroscience*, 22(11):2480–2490, 2010.
57. Jeffrey Mark Siskind. Grounding the lexical semantics of verbs in visual perception using force dynamics and event logic. *J. Artif. Intell. Res. (JAIR)*, 15:31–90, 2001.
58. Kenny Smith, Simon Kirby, and Henry Brighton. Iterated learning: a framework for the emergence of language. *Artificial life*, 9(4):371–386, 2003.
59. Luc Steels. The emergence and evolution of linguistic structure: from lexical to grammatical communication systems. *Connection science*, 17(3-4):213–230, 2005.
60. Luc Steels and Manfred Hild. *Language grounding in robots*. Springer Science & Business Media, 2012.
61. Francesca Stramandinoli, Davide Marocco, and Angelo Cangelosi. The grounding of higher order concepts in action and language: A cognitive robotics model. *Neural Networks*, 32:165–173, 2012.
62. Yuuya Sugita and Jun Tani. Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. *Adaptive Behavior*, 13(1):33–52, 2005.
63. Ilya Sutskever, Oriol Vinyals, and Quoc VV Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
64. Jun Tani. Self-organization and compositionality in cognitive brains: A neurorobotics study. *Proceedings of the IEEE*, 102(4):586–605, 2014.
65. Jun Tani, Masato Ito, and Yuuya Sugita. Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using RNNPB. *Neural Networks*, 17(8):1273–1289, 2004.

66. Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R Walter, Ashis Gopal Banerjee, Seth J Teller, and Nicholas Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *AAAI*, 2011.
67. Michael Tomasello and Michael Jeffrey Farrar. Object permanence and relational words: A lexical training study. *Journal of Child Language*, 13(03):495–505, 1986.
68. David C Van Essen. Functional organization of primate visual cortex. *Cerebral cortex*, 3:259–329, 1985.
69. Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. *arXiv preprint arXiv:1411.4555*, 2014.
70. Anne S Warlaumont, Gert Westermann, Eugene H Buder, and D Kimbrough Oller. Prespeech motor learning in a neural network using reinforcement. *Neural Networks*, 38:64–75, 2013.
71. Juyang Weng, James McClelland, Alex Pentland, Olaf Sporns, Ida Stockman, Mriganka Sur, and Esther Thelen. Autonomous mental development by robots and animals. *Science*, 291(5504):599–600, 2001.
72. Daniel M Wolpert, Zoubin Ghahramani, and Michael I Jordan. An internal model for sensorimotor integration. *Science*, 269(5232):1880, 1995.
73. Yuichi Yamashita and Jun Tani. Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS computational biology*, 4(11):e1000220, 2008.
74. Onur Yürüten, Erol Şahin, and Sinan Kalkan. The learning of adjectives and nouns from affordance and appearance features. *Adaptive Behavior*, 21(6):437–451, 2013.
75. Junpei Zhong. Artificial neural models for feedback pathways for sensorimotor integration. 2015.
76. Junpei Zhong and Lola Canamero. From continuous affective space to continuous expression space: Non-verbal behaviour recognition and generation. In *Development and Learning and Epigenetic Robotics (ICDL-Epirob)*, 2014 Joint IEEE International Conferences on, pages 75–80. IEEE, 2014.
77. Junpei Zhong, Angelo Cangelosi, and Stefan Wermter. Toward a self-organizing pre-symbolic neural model representing sensorimotor primitives. *Frontiers in behavioral neuroscience*, 8, 2014.
78. Junpei Zhong, Cornelius Weber, and Stefan Wermter. Robot trajectory prediction and recognition based on a computational mir-

ror neurons model. In *Artificial Neural Networks and Machine Learning–ICANN 2011*, pages 333–340. Springer, 2011.

79. Junpei Zhong, Cornelius Weber, and Stefan Wermter. Learning features and predictive transformation encoding based on a horizontal product model. In *Artificial Neural Networks and Machine Learning–ICANN 2012*, pages 539–546. Springer, 2012.
80. Junpei Zhong, Cornelius Weber, and Stefan Wermter. A predictive network architecture for a robust and smooth robot docking behavior. *Paladyn, Journal of Behavioral Robotics*, 3(4):172–180, 2012.

Appendix

To calculate the coordinates of the fixed points, we should let $f'(x) = 0$, which means that we need to solve the following equations

$$\begin{cases} -x_1(t) + f(x_3(t)) = 0 \\ -x_2(t) + a \cdot f(x_1(t)) + c \cdot PB = 0 \\ -x_3(t) + b \cdot f(x_2(t)) = 0 \end{cases} \quad (18)$$

The first solution for the first coordinate $[x_1^1, x_2^1, x_3^1]$ is:

$$\begin{cases} x_1^1 = \sqrt{\frac{36N^2 + (6M - 6N^2 + Na)^2}{36N^2b^2 + 36N^2 + (6M - 6N^2 + Na)^2}} \\ x_2^1 = \frac{a}{6} + \frac{M}{N} - N \\ x_3^1 = 6\sqrt{\frac{N^2b^2}{36N^2 + (6M + N(-6N + a))^2}} \end{cases} \quad (19)$$

Similarly, the coordinate of the second fixed point $[x_1^2, x_2^2, x_3^2]$ is calculated by:

$$\begin{cases} x_1^2 = \sqrt{\frac{[2250000(173N-100)^2 + (-30000M + 224727N^2 + 43250Na + 259650N + 25000a + 75000)^2]}{[225000b^2(173N-100)^2 + 225000(173N-100)^2 + (-30000M + 224727N^2 + 43250Na + 259650N + 25000a + 75000)^2]}} \\ x_2^2 = \frac{a}{6} + \frac{-2M}{-1+\sqrt{3}}N - (-\frac{1}{2} + \frac{\sqrt{3}}{2} \cdot N) \\ x_3^2 = 6\sqrt{\frac{b^2 \cdot (1.73N-1.0)^2}{-12M + (1.73N-1.0) \cdot (5.196N+a+3.0)^2 + 36(1.73N-1.0)^2}} \end{cases} \quad (20)$$

And the coordinate of the third fixed point $[x_1^3, x_2^3, x_3^3]$ is given by

$$\begin{cases} x_1^3 = \sqrt{\frac{[2250000(173N+100)^2 + (-30000M + 224727N^2 + 43250Na + 259650N + 25000a + 75000)^2]}{[225000b^2(173N+100)^2 + 225000(173N+100)^2 + (-30000M + 224727N^2 + 43250Na + 259650N + 25000a + 75000)^2]}} \\ x_2^3 = \frac{a}{6} + \frac{-2M}{-1-\sqrt{3}}N - (-\frac{1}{2} - \frac{\sqrt{3}}{2} \cdot N) \\ x_3^3 = 6\sqrt{\frac{b^2 \cdot (1.73N+1.0)^2}{-12M + (1.73N+1.0) \cdot (5.196N+a+3.0)^2 + 36(1.73N+1.0)^2}} \end{cases} \quad (21)$$

For the above solutions, we define the parameters M and N as follows:

$$M = -\frac{a^2}{36} - \frac{b^2}{6} - \frac{1}{3}$$

$$N = \left[-\frac{a^3}{216} + \frac{ab^2}{4} + \frac{-\frac{ab^2}{2} - 1}{12} + \frac{a}{4} + \frac{c \cdot PB}{4} + \sqrt{M^3 + \frac{(-\frac{a^3}{108} + \frac{ab^2}{2} + \frac{-\frac{ab^2}{2} - 1}{6} + \frac{a}{2} + \frac{c \cdot PB}{2})^2}{4}} \right]^{1/3}$$

Although the equations seem to be complicated, remember that variables a , b and c (weights) are constant after training, which means that M is a constant as well. Thus PB value is a first-order variable in the function of N . Similarly, from observation from Eqs.19 - 21 we can see that the solutions are first-order function of variable N , which means that the coordinates of this non-linear system are a first-order function of PB .



Junpei Zhong Junpei 'Joni' Zhong is a junior researcher at Lab for Intelligent Dynamics and Representation, Waseda University, Tokyo. He is also a visiting researcher at Plymouth University, UK. After having a B.Eng from South China University of Technology, an M.Phil degree from the Hong Kong Polytechnic University and a doctor of natural sciences (Dr.rer.nat) from University of Hamburg. Between year 2010 and 2015, he has been engaged in several EU FP7 projects on cognitive robotics

(RobotDoC, ALIZ-E and POETI-CON++) in Germany and the UK.

His research is committed to build learning models for cognitive robotics. In most cases, such models are inspired by neuroscience and cognitive science.



Martin Peniak Martin Peniak is a development engineer at Cortexica Vision Systems, London, UK, after he graduated from Plymouth University, UK. His research interests include parallel programming, cognitive robotics, astrophotography.



Jun Tani Jun Tani received a BS in Mechanical Engineering from Waseda University, a dual MS in Electrical Engineering and Mechanical Engineering from the University of Michigan (Ann arbor), and a Dr. Eng from Sophia University in Tokyo. He started his research career in Sony Laboratory in 1990. He had been a team leader of the Laboratory for Behaviour and Dynamic Cognition, Brain Science Institute, RIKEN in Tokyo for 12 years until May 2012. He also held the position of visiting associate

professor at the University of Tokyo between 1997 and 2002. Now he is a

full professor in Electrical Engineering Department in KAIST in Korea. His research interests include neuroscience, psychology, phenomenology, complex adaptive systems, and robotics.



Tetsuya Ogata Tetsuya Ogata received the BS, MS and DE degrees in Mechanical Engineering, in 1993, 1995 and 2000, respectively, from Waseda University. He was a research fellow of JSPS, a research associate of Waseda University, a research scientist of RIKEN Brain Science Institute, and an associate professor of Kyoto University. He is currently a professor of Faculty of Science and Engineering, Waseda University. Since 2009, he has been a JST PRESTO researcher. His research interests

include human-robot interaction, dynamics of human-robot

mutual adaptation and inter-sensory translation in robot systems. He is a member of IEEE, RSJ, JSAI, IPSJ, JSME, SICE, etc.



Angelo Cangelosi Angelo Cangelosi is Professor of Artificial Intelligence and Cognition and the Director of the Centre for Robotics and Neural Systems at Plymouth University (UK). He studied psychology and cognitive science at the Universities of Rome La Sapienza and at the University of Genoa, and has been visiting scholar at the University of California San Diego and the University of Southampton. Cangelosi's main research expertise is on language and cognitive modelling in humanoid robots, on

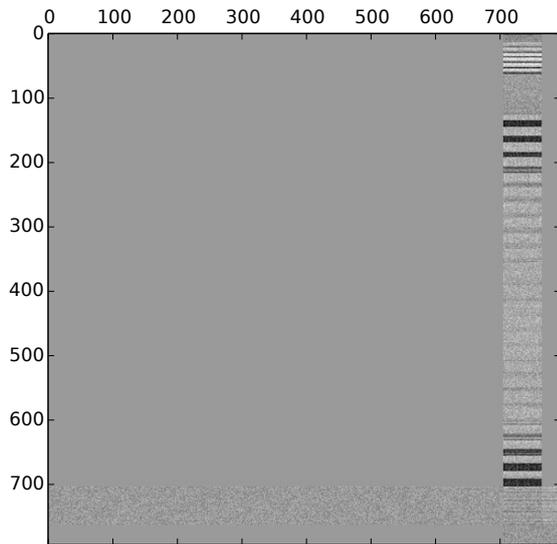
language evolution in multi-agent systems, and the application of bio-

inspired techniques to robot control (e.g. swarm of UAVs). He is the

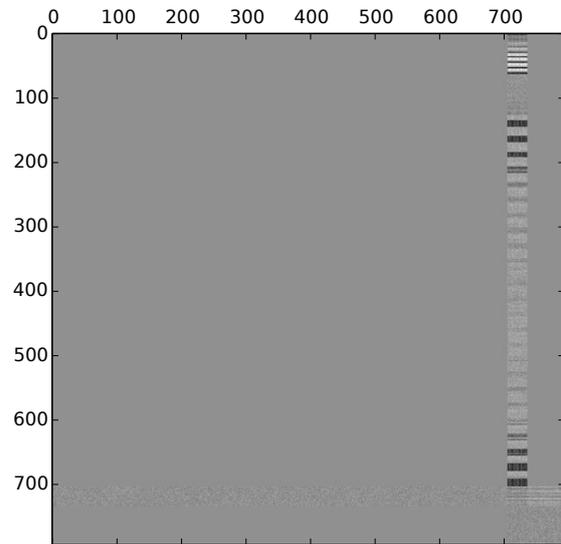
coordinator of the Marie Curie ITN "RobotDoC: Robotics for Development of Cognition" (2009-2014) and the UK EPSRC project "BA-BEL: Bio-inspired Architecture for Brain Embodied Language" (2012-2016), and of the FP7 project "ITALK" completed in 2012. Cangelosi has produced more than 200 scientific publications, and has chaired numerous workshops and conferences including the IEEE ICDL-EpiRob 2011 Conference (Frankfurt, August 2011). In 2012-13 he was Chair of the International IEEE Technical Committee on Autonomous Mental Development. He will be Editor-in-Chief of the IEEE Transactions from Autonomous Development from January 2015, and currently is Editor of the journal Interaction Studies. He is co-author, with Matt Schlesinger, of the forthcoming book "Developmental Robotics: From Babies to Robots" (MIT Press, 2015).

Fig. 6: Weight Visualization by Input Removal

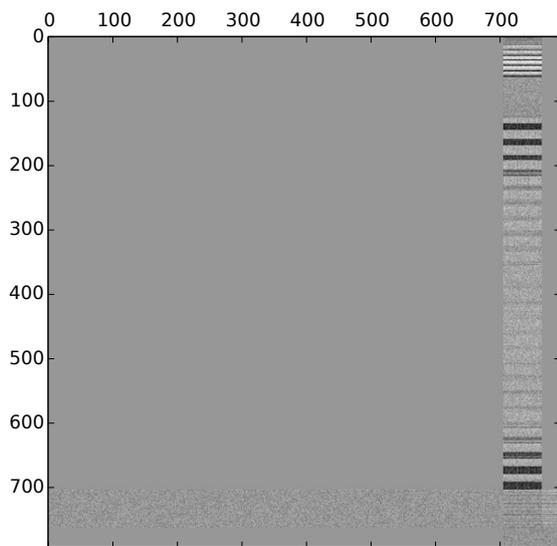
(a) Weight Matrix of Normal Training (base-line)



(b) Weight Matrix without Verb Input



(c) Weight Matrix without Noun Input



(d) Weight Matrix without Visual Input

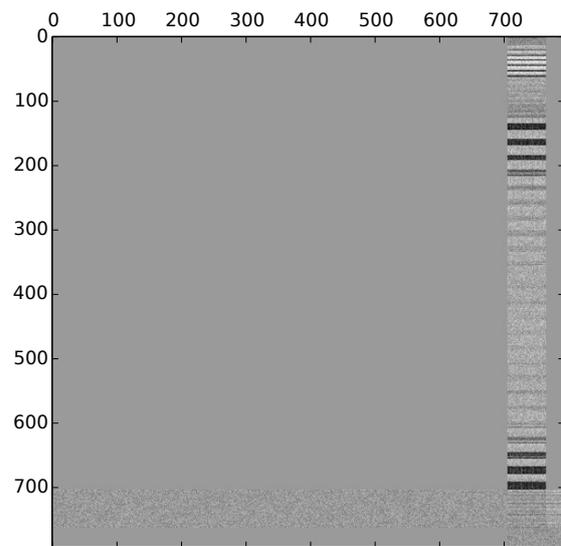


Fig. 7: Principle Component Analysis on the C_f neurons

Comparison of neural activation (after PCA) on the C_f layer shows that the sequences with different nouns are clustered closer than those with different verbs. Particularly we can compare combinations of $[0, 0.0 - 0.3]$ and $[0.1, 0.0 - 0.2]$.

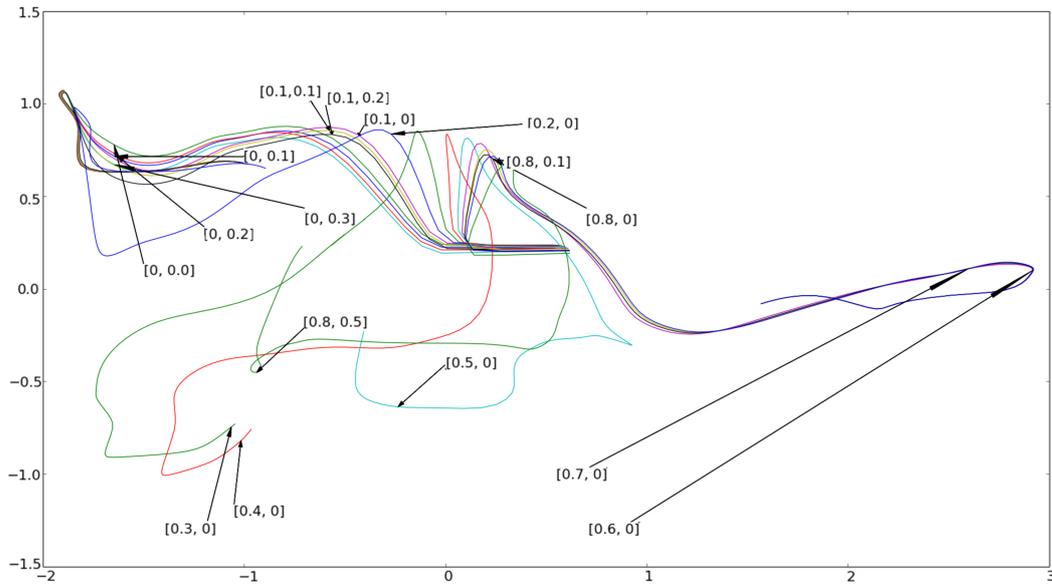


Fig. 8: Principle Component Analysis on the C_f neurons (with different nouns): Comparison of PCA processed neural activation shows that the sequences with different nouns differ at the beginning and at the middle of the trajectories.

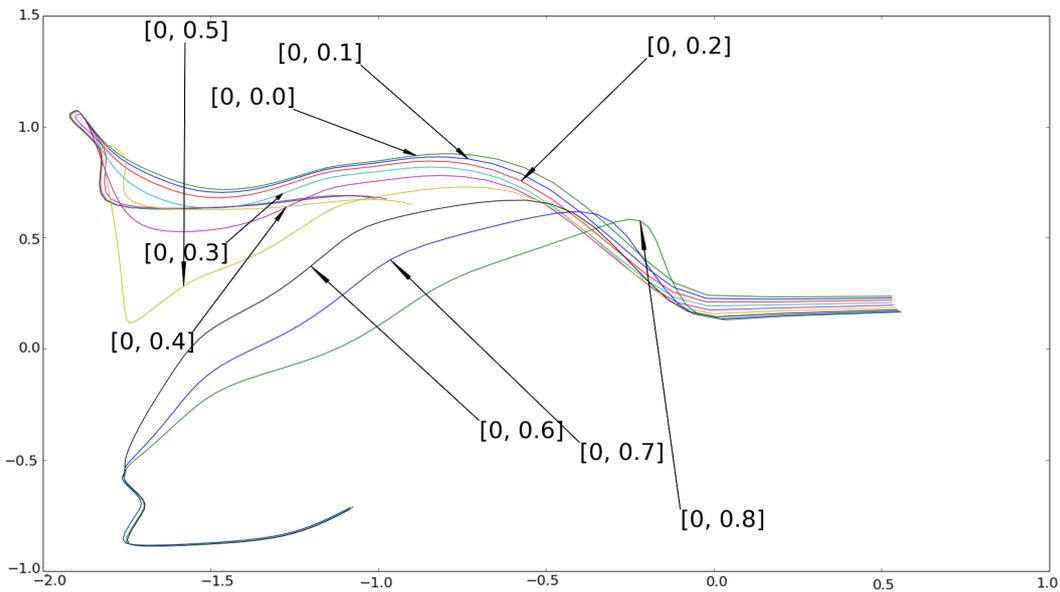


Fig. 9: Principle Component Analysis on the C_f neurons (with different object locations): Comparison of PCA processed neural activation shows the sequences with different visual inputs result in very little divergence in the trajectories, which mainly occurs in the middle of the trajectories.

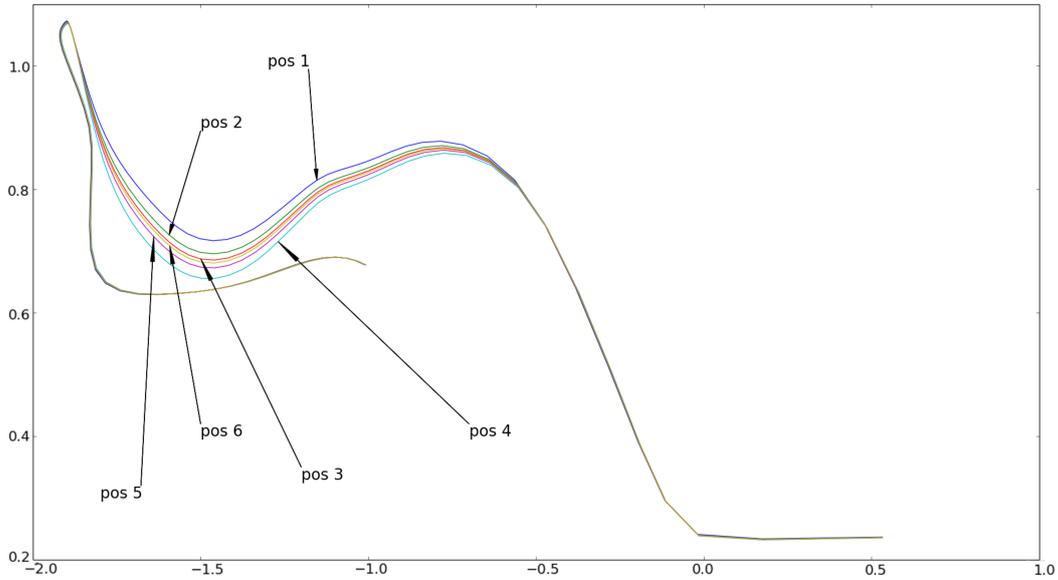


Fig. 10: Principle Component Analysis on the C_s neurons

Comparison of PCA processed neural activation shows that the differences of the sequences are mainly present in the verbs on the C_s layer.

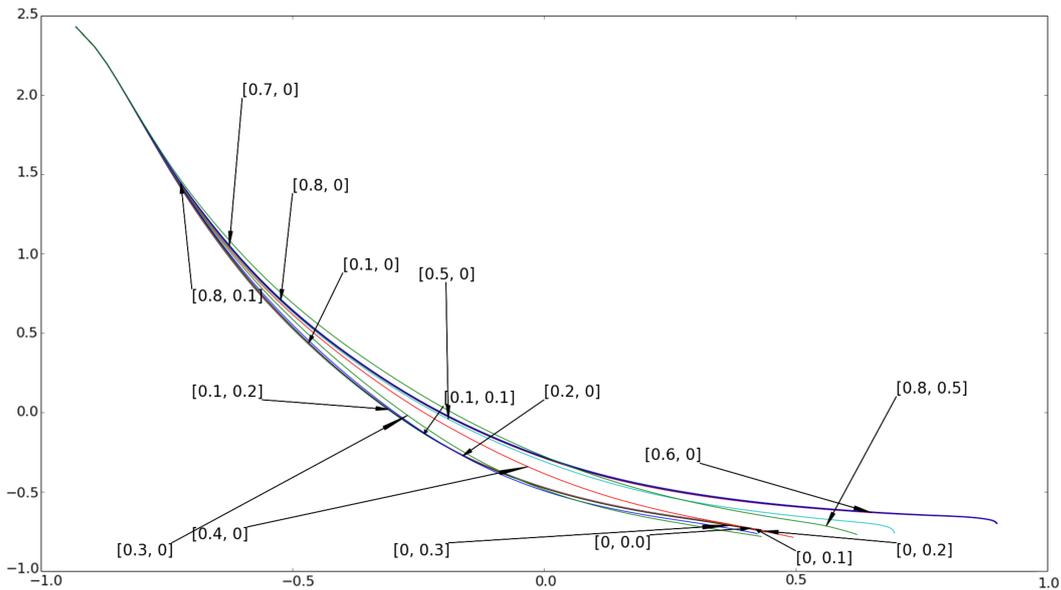


Fig. 11: Principle Component Analysis on the C_s neurons (with different nouns)

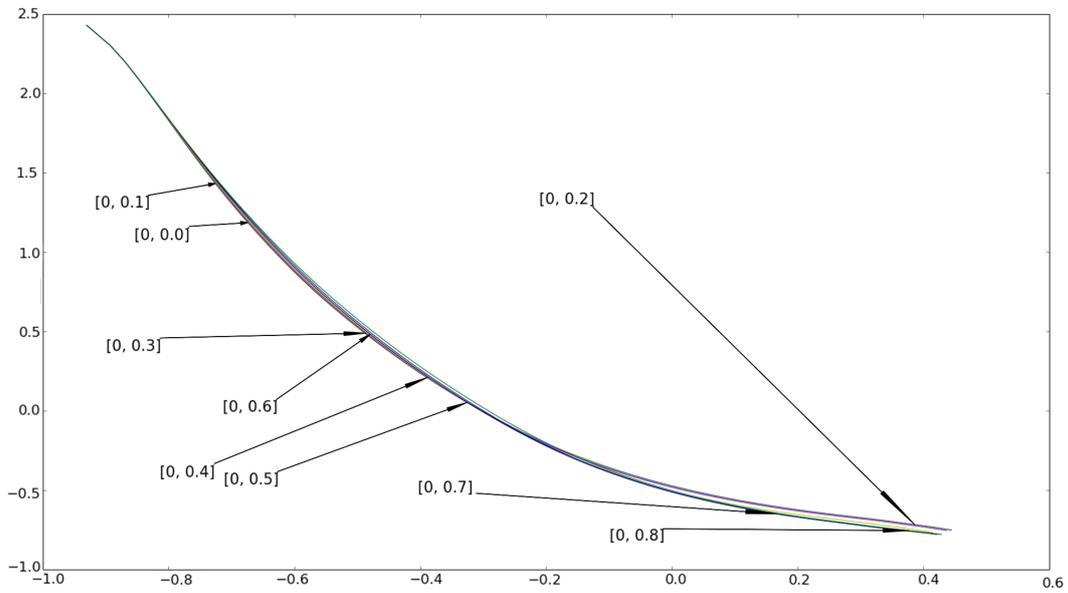


Fig. 12: Principle Component Analysis on the C_s neurons (with different object locations)

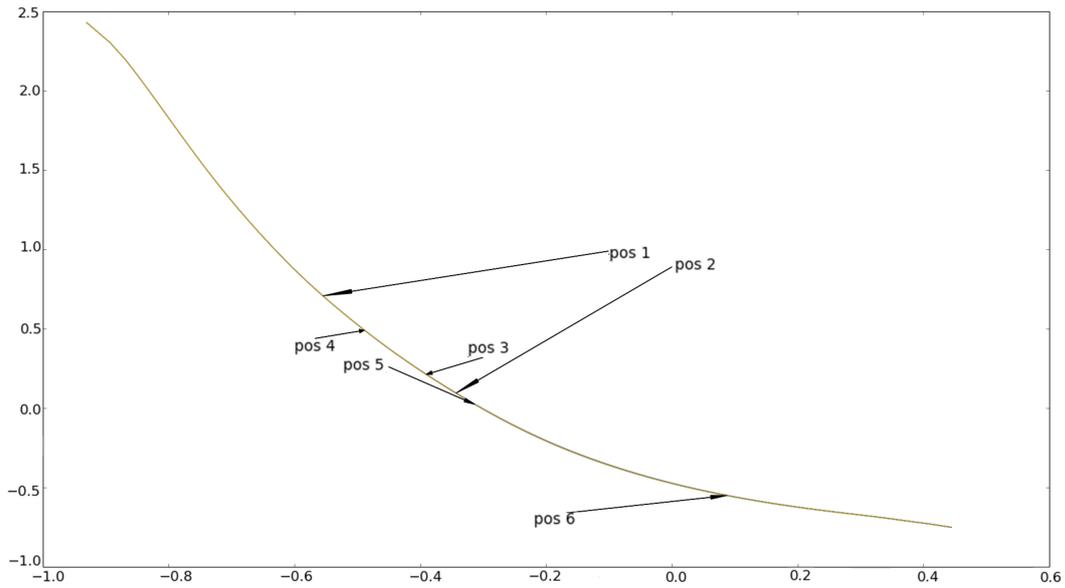


Fig. 13: A Simple Recurrent Network with Parametric Bias Units

