W. Erlhagen, E. Bicho, "*Dynamic Field Theory (DFT): Applications in Cognitive Science and Robotics*", 10 pages, **euCognition: The European Network for Advancement of Artificial Cognitive Systems**, Network Action 057-1, 9 Janeiro, 2009. (http://www.eucognition.org/euCognition_2006-2008/white_papers.htm)

Dynamic Field Theory (DFT): Applications in Cognitive Science and Robotics

Wolfram Erlhagen[1] and Estela Bicho[2]
[1]Dep. of Mathematics for S&T, [2]Dep. of Industrial Electronics
University of Minho, Portugal
(wolfram.erlhagen@mct.uminho.pt, estela.bicho@dei.uminho.pt)

Introduction

In recent years, there has been an increasing interest in the dynamical systems approach to cognitive science that challenges the more traditional, computational view on cognition ([1], [2], [3]). The heart of the dynamic approach is the hypothesis that natural cognition embedded in a nervous system and in an environment unfolds continuously in time as the result of the complex interplay between a variety of internal and external forces. Proponents of this new direction argue that cognitive processes can be conceptualized in terms of position, distance and trajectories in the space of possible states of the system. Consequently, the primary focus of their theoretical and empirical research program is on understanding the nature of change in time. The dynamic hypothesis strongly contrasts with the computational approach that explains cognitive behaviour as a rule-based manipulation of static internal representations.

The theory of nonlinear dynamical systems is the mathematical framework that offers essential concepts and tools such as *attractor*, *repeller*, *stability*, *bifurcation*, *trajectory in state space* and the like. However, these concepts are certainly necessary but not sufficient for a dynamical systems approach to higher cognitive capacities. The problem is that flexible and intelligent behaviour very often depends on past and future events of the environment that are not represented in the immediate stream of sensory signals. The forces that are supposed to shape the unfolding trajectories are thus not available all time. What is needed is the notion of complex persisting inner states that allow the cognitive agent, for instance, to compensate for temporally missing sensory input, or to anticipate future environmental inputs that may inform the decision about a specific goal-directed behaviour [1].

Dynamic Field Theory or DFT provides a theoretical framework that bridges the representational gap in the dynamic system approach ([4], [5]). Dynamic fields formalized by differential equations have been originally introduced by Wilson and Cowan [6] and Amari [7] in the 1970's as simplified mathematical models for pattern formation in neural tissue. The architecture of this model family reflects the hypothesis that strong recurrent interactions in local populations of neurons form a basic mechanism for cortical information processing. This hypothesis is supported by anatomical and psychological findings showing that a cortical neuron gets its largest input from neighbouring cells. The recurrent interactions cause non-trivial dynamic behaviour in neural populations. Most importantly, population activity which is initiated by time-dependent external signals may become self-sustained in the absence of any external input. Such attractor states of the population dynamics may be used by the nervous systems to guide complex behaviour that goes beyond a simple input-output mapping. On this view, DFT complements the dynamic approach by adding the notion of sub-symbolic dynamic representations that are consistent with fundamental principles of cortical information processing.

In the following we briefly discuss representative examples that shall illustrate the basic concepts of DFT and its application in different cognitive domains (for a recent survey of the mathematical treatment see [5]). As embodied validations of the approach, we include studies in which autonomous robots were endowed with some cognitive capacities such as decision making, memory, action understanding or anticipation. The reader is encouraged to visit the original publications for proper treatment.

Basic concepts

Dynamic fields are defined over metric dimensions that in most applications reflect parameters of the experimental setup such as for instance visual space, object features or movement parameters. A localized peak of activation within the field represents metric information about the underlying dimension. In analogy to the concept of neural population coding, the localized peak may be interpreted as the activation pattern of a pool of neighbouring neurons with similar tuning properties in an otherwise inactive population [8]. The peak evolves under the influence of time-dependent external input but is mainly shaped by the recurrent interactions within the population. Due to the interactions the peak is self-sustained and may thus serve to memorize information. However, in order to perform a working memory the field dynamics has to be bi-stable. The attractor state of a persistent firing pattern coexists with a stable homogenous activation distribution that represents the absence of information. A sufficiently strong transient input specifying a particular instance of the metric dimension is able to "switch on" the memory function by destabilizing the uniform rest state. A suitable inhibitory signal, on the other hand, may destabilize the excited state, resulting in a decay of activation back to resting state. The critical parameter that decides whether or not the uniform state becomes unstable is input strength. Below a certain threshold, inhomogeneous solutions of the field dynamics exist that are mainly input driven. In many applications this regime is exploited to represent for instance ambiguous prior information about the dimension spanned by the field. Much like prior distributions in the Bayesian sense, multi-modal patterns of subthreshold activation may encode the probability of choices [9].

Within the DFT framework, the existence of a memory function is closely linked to decision making. A self-sustained activation peak is the result of the interplay between excitatory and inhibitory forces within the field. Excitatory interactions between neighbouring neurons are counterbalance by lateral inhibition. Since inhibitory interactions dominate at larger distances, a competition process takes place whenever two or more suprathreshold inputs are simultaneously applied. At the field site that integrates the highest level of sensory or other evidence a single activation peak will emerge. It suppresses the activation elsewhere in the field below resting level. Once the stable attractor state has been reached, the decision represented by the peak location is protected against noise and fluctuations in the input stream. This does not mean, however, that the peak is completely immune against change. Localized input that spatially overlaps with the peak position causes the peak to drift in the direction of this input. Mathematically, the attractor can be described as being continuous along the dimension spanned by the field. This property of the field dynamics has been exploited for instance to explain and apply the capacity of tracking time-varying input in real time ([10], [11]).

Although the formation of spatially localized peaks is certainly the backbone of most DFT-based models thus fare, it is important to note that the field dynamics - if properly tuned - may exhibit other types of self-organizing patterns (e.g., .travelling pulse and wave, active transient etc., see [12] for a recent overview) that play an important role in some applications of the dynamic field approach to cognition (see the examples below).

Applications in Cognitive Science: Action, Perception and Reasoning

The dynamic field approach has been first developed and applied in the context of simple sensorimotor decisions ([13], [14], [15]). Timed-movement initiation paradigms reveal that plans for goal-directed arm or eye movements evolve continuously in time. Early during action preparation an externally forced decision to move reflects the prior knowledge about the task (e.g., number and probability of possible targets) whereas later in time the decision is dominated by the visual information specifying the target. In the DFT model [14], a localised activation peak in the movement parameter field evolves in time under the influence of the two inputs sources represented by activation patterns in two connected fields. Due to the constant input from the task field, the build up starts from a pre-activated initial state reflecting prior task knowledge. Since the level of pre-activation affects the rate at which the activation peak rises, all classical effects of task environment on reaction time can be modelled. Moreover, the field model also accounts for experimental findings showing that not only the informational content (e.g., number and probability of choices) but also the metric of the movement parameter space matters. A simple example is a two choice task in which the relative distance between the possible directions of the movement is systematically varied. A general experimental finding is that reaction time decreases with decreasing distance. In the model, the two-choice setting lead to a bimodal preshape. For metrically close directions, however, the two preshape inputs start to overlap, resulting in a higher level of pre-activation. As a consequence, the buildup of the peak in response to a specific input is faster compared to the case of two movement directions metrically fare from each other.

It is worth mentioning that the dynamic field concepts can be generalized to the case of multiple movement parameters by introducing multidimensional activation fields. A two-dimensional field spanned over movement direction and movement amplitude has been used for instance to explain differences in the process of specification of the two parameters as revealed in precuing paradigms [14].

The DFT approach to sensorimotor decisions has been extended to address fundamental questions in developmental cognitive science. An excellent example is the DFT model of Piaget's classic 'A-not-B' error [16]. In this task, an infant is faced with two similar opaque containers with lids. The experimenter hides a toy under the lid at location A and the infant is trained to reach for the toy. If after several A trials the object is hidden for the first time at location B and a short delay is imposed, younger infants nevertheless reach toward container A. This typical 'A-not-B' error is absent in older infants. The DFT model explains this change in behaviour as a shift from an input-driven regime of decision making in younger infants to an interaction-dominated regime in older children (but see [17] for a recent refinement of the model). The motor decision field is spanned over the space of reachable directions. It evolves under the influence of the specific input representing the location of the toy, input from the task field representing the perceptual layout of the work space, and input from the preshape field representing the memory trace of previous A-reaches. If the decision field is input-driven, that is, the transient input representing the toy at B does not reach the threshold for triggering the cooperative forces within the population, the activation pattern at the respective field site will decay during the delay towards resting level. As a result, movement history represented by the preshape field will start to dominate movement selection, leading to the ‘A-not-B’ error. If, by contrast, the field is able to self-stabilize activation patterns, the evolving peak at the B location will suppress the conflicting input from the preshape field.

During the delay period between the presentation of the cue at a certain location and movement onset, the peak serves as spatial working memory. However, since the peak attractor is continuous, spatially overlapping preshape input will cause a continuous drift of the localized activity pattern whenever there exists no additional input that may ‘fix’ the peak at the cued position. The predicted distortion over time of spatial memory has been indeed observed in experiments with goal-directed actions to remembered locations with no salient location cues [18].

Cognitive vision is another attractive area for DFT research [19]. The basic concept of self-stabilized internal representations provide the means to cope with some of the fundamental problems any cognitive agents has to solve in order to organize goal-directed behaviour in a continuously changing world. Very importantly, the system has to compensate for long and unreliable delays involved in the processing of visual information about moving objects. As revealed in psychophysical experiments, the position percept very often does not agree with the physical measurements. Furthermore, a continuous stream of sensory signals that may guide behaviour is not guaranteed in cluttered environments. There is thus a need for mechanisms that allow the system to extrapolate past trajectory information into the future.

The DFT-based model of object location in visual space has been originally developed to explain nonlinear interaction effects found in neuronal population activity of primary visual cortex [20]. It has been later extended to bridge the gap to the perceptual level. The model accounts for various systematic localization errors observed with moving and stationary targets ([21], [22], [23], [24]). In response to an object in real or apparent motion, the field develops a self-stabilized travelling pulse of activation that locks to the stimulus. Its peak position is taken as a neural correlate for the location percept. Due to the lateral interactions within the population, neurons encoding future stimulus positions become pre-activated, leading to a path-dependent facilitation of processing. In the model, the latency of the suprathreshold population response to the external input may be further reduced by integrating additional preshape input from other information sources such as for instance actions plans directed toward the moving object. Moreover, like a physical object in motion, the internal representation of target position cannot be halted instantaneously upon stimulus offset but continues to travel for some time. In line with experimental findings in position judgement tasks, the model predicts a perceptual overshooting at the end of the motion trajectory whenever the cooperative forces within the population are sufficiently strong. The extent to which the internal representation extrapolates past trajectory information depends on a single model parameter that controls the interplay between excitation and inhibition within the field. If properly tuned, the self-organizing wave may fill in missing trajectory information whenever the sensory input is temporally disrupted due to occluding surfaces for instance [25].

Decision making, memory and object permanence are all examples of basic cognitive functionalities. Recently, attempts have been made to extend the DFT approach to higher cognitive capacities such as off-line reasoning and prediction that are often crucially involved in the control and coordination of real action in real time. To give a concrete example from the domain of social cognition, a fluent and efficient coordination of actions and decisions among partners in a joint action task requires that each individual is able to understand others' action goals. The DFT model of goal inference [26] implements the idea that action simulation supports action understanding. Without the need to rely on any explicit symbolic communication, we understand the purpose of observed actions by internally replicating action effects using our own motor repertoire. The model has been validated in tasks which involve goal-directed reaching-grasping-placing sequences. The model architecture consists of various reciprocally coupled dynamic fields that represent in their firing patterns action means, action goals and contextual cues. In the action simulation layer (ASL) learned action chains composed of populations encoding different motor acts exist that are linked to specific placing goals or end states. Individual chains may be preshaped by connected populations representing contextual cues and prior task information. Through mappings between populations encoding congruent motor acts in the action observation (AOL) and the action simulation layer, the chain is triggered whenever input from AOL indicates the observation of a specific motor act. The self-stabilizing properties of the neural populations forming the chain ensure that the associated goal representation reaches a suprathreshold activation level even in cases where only partial visual information about the action performed by another person is available (e.g., only the reaching toward the object is observed). In a joint action context, the information about the inferred goal may then be used by the observer to select an appropriate complementary action sequence.

It is worth stressing that the dynamic field model within which the high-level reasoning capacity is realized represents a distributed but fully integrated dynamical system that emerges from sensorimotor origins. The same action chains that support overt motor behaviour are covertly used during action observation to make sense of the behaviours of others. The populations forming the chains encode entire goal-directed motor primitives such as grasping and placing that abstract from the fine details of the hand movements. For such a high-level motor vocabulary the metric structure of the underlying space, which is a core concept of DFT, is not directly observable. However, the metric may still be defined operationally by the degree of overlap of neuronal representations. In the present implementations, a motor acts such as grasping is represented by a localized activation pattern in an otherwise inactive population. The pool of grasping neurons that becomes activated above threshold gets input from connected populations. In a modelling attempt to explain the development of goal-directed chains, it has been recently shown that the interplay of the field dynamics and a Hebbian learning dynamics may result in the emergence of distinct, non-overlapping subpopulations of grasping neurons that become linked during learning to specific goal representations [27]. In other words, depending on the ultimate goal of the action sequence in which the grasping behaviour is embedded, the localized activity pattern in the grasping field will evolve at different locations of the field. On this view, the similarity on the goal level defines the metric of the field, that is, the extent to which subpopulations will overlap. The development of such goal-directed representations has important implications for cognitive behaviour since it allows individuals to act at the time of the grasping in anticipation of others' motor intentions.

Applications in Robotics

An active area of research is the validation of the neuro-plausible mechanisms and concepts of the DFT approach to cognition in autonomous artificial agents. Implementations on robots controlled in closed-loop by sensory information provide feedback about hidden assumptions and missing links to the sensory and motor surfaces. A major challenge is to demonstrate that inner persistent states and internal simulations are not a costly information processing bottleneck [28] but actually support fluent real-world behaviour.

DFT has been first introduced into the domain of robotics in navigation tasks with the goal to modify the behaviour of an autonomous vehicle through memorized sensory information ([29], see [30], [11], [31] for follow up studies). In the DFT-based control architectures, the fields are spanned over the heading direction of the robot relative to an arbitrary but fixed frame of reference. The path planning field integrates on-line information about the target location and the location of obstacles coming form on-board sensors. In addition, it receives input from an object memory field that contains self-stabilized activation peaks. They represent the location of previously detected obstacles that are currently invisible to the sensors. This memorized information is coupled into the movement planning field in the form of inhibitory input, defining undesirable heading directions. Since the field dynamics supports the existence of a single localized activation peak, decision can be made and stabilized even in navigation tasks with several potential targets. The demonstrations on vehicles equipped with low-level sensors and controlled by non-linear attractor dynamics demonstrate that the navigation behaviour in cluttered environments is quite smooth and robust [11]. The robustness is a direct consequence of the strong recurrent interactions within the fields that amplify and stabilize noisy and corrupted sensory inputs.

Smooth real-time tracking of a moving target (e.g., another agent) is possible by exploiting that the self-stabilized activation peak in the target layer of the control architecture follows the continuously moving sensory input. This sensor-dependent tracking process will be disrupted, however, if the target disappears behind an occluding surface. The solution of this problem offered by DFT is a self-stabilized wave which travels in the direction of the moving target with a speed that can be controlled depending on the task. For instance, to anticipate future locations of the moving target the speed of the travelling wave may be chosen slightly faster than the actual velocity of the target. This predictive mechanism has been applied in a joint action task in which two agents (robot-robot or robot-human teams) have to search objects in a common workspace with the goal to transport them to a predefined area ([5], [32]). To guarantee an efficient division of the search space without the help of direct communication, each teammate has to infer from the partner's motion future positions and possible action goals. The self-stabilized wave representing the predicted motion trajectory couples into the decision field of the observing robot as inhibitory input. As a consequence, the decision about the object toward which the robot will move appears to be biased by sensed or memorized targets in other locations of the workspace.

The neuro-cognitive mechanisms implemented by the dynamic field approach to cognition represent a promising research direction for the domain of human-robot interaction (HRI). The basic idea followed by a number of studies ([33],[34],[35],[38]) is that in order to advance towards socially aware robots able to cooperate with humans in a natural and efficient manner, the robot design should reflect similar processing principles.

An example on the perceptual level is the interactive learning of labels for different objects that are manipulated in joint action tasks [36]. The dynamic field architecture consists of various coupled two-dimensional feature label fields that get input from the camera system. Object features like colour, shape and size represent one dimension of a field whereas the object label is represented along the second dimension. The output of all label feature fields is integrated in a decision layer that contains one node for each label. There exist lateral connections between the feature label fields that ensure that during the recognition process the same labels belonging to different features fields excite each other. Feedback connections from the output field, on the other hand, guarantee that activation peaks in the feature label fields that do not match the decision in the output field are extinguished. During the supervised training phase, a correct label is given to the robot if recognition fails. Learning within the distributed network takes place in form of a buildup of preshape (or memory trace) at positions where a peak has emerged during successful recognition trials.

The DFT model of action understanding has been integrated in control architectures to endow robots with the capacity of goal-directed imitation ([5], [37]). It has been validated in the context of grasping-placing sequences. The defining feature of the three-layer field architecture is a learned mapping from motor primitives in the action observation layer onto motor primitives in the action execution layer. It implements a matching on the goal and not on the movement level. The architecture can thus cope with differences in embodiment, task constraints and motor skills between teacher and robot. A self-stabilized activation peak representing a certain goal biases the decision process in the action execution layer toward the selection of motor primitives that can be used by the robot to achieve that goal. Moreover, the existence of goal-directed action chains in the middle layer of the field architecture allows the robot to act on a goal that it has to infer since the end state of the sequence was hidden from view or the user accidently failed to place the object at the desired location. If the meaning of an action is understood, further knowledge may be transferred from an experienced teacher to the robot in imitation paradigms. For instance, knowledge where o place a certain class of objects (e.g., defined by colour or shape) may be acquired and will be represented in learned synaptic links between neural populations encoding the goal and specific object features.

For cooperative joint tasks, imitating the human behaviour is normally not beneficial for the robot and the team. The interpretation of observed actions in terms of goals should lead to the selection of an appropriate complementary behaviour. Both processes, goal inference and the action selection, depend on the situational context. A dynamic field architecture that establishes a context dependent mapping from observed actions onto complementary actions has been recently developed and tested in a task in which a robot and a human jointly construct a toy robot [38]. The state of the construction that defines which sub-goals are currently available for the team is represented by neural populations in the common sub-goal layer (CSGL) of the field architecture. If the input from this layer and from the intention layer (IL) representing the inferred goal of the partner converge on the representation of a particular action sequence in the action execution layer (AEL), this action sequence will normally win the competition process between all possible actions. If a mismatch between the inferred goal and the possible sub-goals exist, the robot is able to detect this error since an activation peak in the error monitoring layer (EML) will automatically appear. A selection of communicative gestures like pointing to ‘explain’ the error to the human partner may then constitute an appropriate complementary behaviour. Interestingly, how social the robot behaves can be controlled by a simple adaptation of the synaptic weights form the intention layer to the action selection layer. If the weights are relatively weak, the action selection process of the robot will be dominated by information sources other than the behaviour of the partner.

References

[1] A. Clark, The dynamical challenge, *Cognitive Science: A Multidisciplinary Journal* 21: 461-481, 1997

[2] T. van Gelder, The dynamical hypothesis in cognitive science, *Behav. Brain Sci*. 21: 615-665, 1998

[3] S. Schöner, Dynamical Systems Approaches to Cognition, In: The Cambridge Handbook of Computational Psychology, R. Sun (ed.), pp. 101-125, Cambrige University Press, 2008

[4] J. P. Spencer, G. Schöner, Bridging the representational gap in the dynamic systems approach to development, *Developmental Science*, 6, 392-412, 2003.

[5] W, Erlhagen and E. Bicho, The dynamic neural field approach to cognitive robotics, *Journal of Neural Engineering* 3:R36-R54, 2006

[6] H.R. Wilson, J. D. Cowan, A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue *Kybernetik (Biol. Cybern.)* **13** 55–80, 1973

[7] S. Amari, Dynamics of pattern formation in lateral-inhibitory type neural fields. *Biological Cybernetics* 27: 77–87, 1977.

[8] W. Erlhagen, A. Bastian, D. Jancke, A. Riehle and G. Schöner, The distribution of population activation (DPA) as a tool to study interaction and integration in cortical representations, *Journal of Neuroscience Methods*, 94:53-66, 1999.

[9] R. H. Cuijpers, W. Erlhagen, Implementing Bayes’ Rules with Neural Fields, In: ICANN 2008, Part II, *LNCS 5164*, V. Kurková et al. (eds.), pp 228-237, Springer-Verlag, 2008

[10] S. Wu, S. Amari, Computing with Continuous Attractors: Stability and Online Aspects, *Neural Computation*, 17:2215-2239, 2005

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

[12] S. Coombes, Waves, bumps, and patterns in neural field theories, *Biological Cybernetics*, 93: 91-108, 2005

[13] G. Schöner, K. Kopecz, W. Erlhagen, The dynamic neural field theory of motor programming: Arm and eye movement”. In P. Morasso, V. Sanguineti, (eds.), *Self-Organization, Computational Maps and Motor Control*, Elsevier Science, pp. 271—310, 1997.

[14] W. Erlhagen, G. Schöner, Dynamic field theory of motor preparation, *Psychological Review*, 109:545-572, 2002.

[15] C. Wilimzig, S. Schneider, G. Schöner, The time course of saccadic decision making; Dynamic field theory, *Neural Networks*, 19:1059-1074, 2006

[16] E. Thelen, G. Schöner, C. Scheier, L. Smith, The dynamics of embodiment: A field thoery of infant preservative, *Brain and Behavioural Sciences* 24; 1-33, 2001

[17] G. Schöner, E. Dineva, Dynamic instabilities as mechanisms for emergence, *Developmental Science*, 10: 69-74, 2007

[18] A. R. Schutte, J. P. Spencer, G. Schöner, Testing the dynamic field theory: Working memory for locations becomes more spatially precise over development. *Child Development*, 74, 1393-1417.

[19] M. Giese, Dynamic neural field theory of motion perception, Kluwer Academic Publishers, Boston, 1998

[20] D. Jancke, W. Erlhagen, H. Dinse, A. Akhavan, M. Giese, A. Steinhage and G. Schöner, Parametric population representation of retinal location: Neuronal interaction dynamics in cat primary visual cortex. *Journal of Neuroscience,* 19:9016-9028, 1999.

[21] W. Erlhagen and D. Jancke, The role of action plans and cognitive factors in motion extrapolation: A modelling study, *Visual Cognition* 11: 315-341, 2004.

[22] D. Jancke, W. Erlhagen, G. Schöner, H. Dinse, Shorter latencies for motion trajectories than for flashes in population responses of cat primary visual cortex, *Journal of Physiology* 556.3, 971-982, 2004.

[23] D. Jancke, W. Erlhagen, Bridging the gap: A model of common neural mechanisms underlying the Fröhlich effect, the flash-lag effect, and the representational momentum effect, In: Problems of Space and Time in Perception and Action, R. Nijhawan and B. Khurana (eds.), Cambridge University Press , in press

[24] D. Bocianski; J. Müsseler, W. Erlhagen, Relative mislocalization of successively presented stimuli, *Vision Research*, 48: 2204-2212, 2008

[25] W. Erlhagen, Internal Models for Visual Perception, *Biological Cybernetics* 88: 408-417, 2003

[26] W. Erlhagen, A. Mukovsky , E. Bicho, A dynamic model for action understanding and goal-directed imitation, *Brain Research,* 1083:174-188, 2006

[27] W. Erlhagen, A. Mukovskiy, F. Chersi, E. Bicho, On the Development of Intention Understanding for Joint Action Tasks, Proceedings of the 6th IEEE International Conference on Development and Learning, Imperial College London, 11-13 July 2007 (ISBN: 1-4244-1116-5).

[28] R. Brooks, Intelligence without representation, *Artificial Intelligence,* 47:139-159, 1991

[29] G. Schöner, M. Dose, C. Engels, Dynamics of behavior: theory and applications for autonomous robot architectures *Robot. Auton. Syst.* 16: 213–45, 1995

[30] C. Engels, G. Schöner, Dynamic fields endow behavior-based robots with representations, *Robot. Auton. Syst.*., 14:55-77, 1995

[31] M. Quoy, S. Moga, P. Gaussier, Dynamical neural networks for planning and low-level robot control, *IEEE Trans. Syst. Man Cybern.* A 33: 523–32, 2003

[32] N. Hipólito, L. Louro, E. Bicho, W. Erlhagen, A neuro-inspired architecture for goal inference and decision making in joint action tasks, Proceedings of the workshop

on Robotics and Mathematics, RoboMat 2007, International Center for Mathematics (CIM), University of Coimbra, Portugal, 2007

[33] R. Menzner, A. Steinhage, W. Erlhagen, An Architecture for Interactive Robots based on Nonlinear Dynamics, In: Springer Tracts in Advanced Robotics, Vol. 14, Advances in Human-Robot Interaction, pp. 301-320, Prassler, E.; Lawitzky, G.; Stopp, A.; Grunwald, G.; Hägele, M.; Dillmann, R.; Iossifidis, I. (Eds.) , Springer Verlag, 2005 (ISBN: 978-3-540-23211-7)

[34] T. Bergener, C. Bruckhoff , P. Dahm, H. Janssen, F. Joublin, R. Menzner, A. Steinhage, W. van Seelen W, Complex behavior by means of dynamical systems for an anthropomorphic robot, *Neural Networks*, 12: 1087–99, 1999

[35] W. Erlhagen**,** A. Mukovskiy, E. Bicho, H. Bekkering**,** Development of action understanding in imitation tasks, Proceedings (CD) of the workshop on Social Mechanisms of Robot Programming by Demonstration, 2005 *IEEE International Conference on Robotics and Automation* , ICRA 2005, Barcelona, Spain, 2005

[36] C. Faubel, G. Schöner, Learning to recognize objects on the fly: A neurally based dynamic field approach. *Neural Networks*, 21:562-576, 2008

[37] W. Erlhagen, A. Mukovskiy, E. Bicho, G. Panin, C. Kiss, A. Knoll, H. van Schie, H. Bekkering, Goal-directed Imitation in Robots: a Bio-inspired Approach to Action Understanding and Skill learning, *Robot. Auton. Syst.* 54(5): 353-360, 2006

[38] E. Bicho, L. Louro , N. Hipólito, W. Erlhagen, A dynamic neural field architecture for flexible and fluent human-robot interaction, Proceedings of the 2008 International Conference on Cognitive Systems, pp. 179-185, University of Karlsruhe, Germany, April 2008