Modeling Motivation in MicroPsi 2

Joscha Bach

Massachusetts Institute of Technology, Cambridge, MA
joscha@mit.edu

Abstract. The MicroPsi architecture combines neuro-symbolic representations with autonomous decision making and motivation based learning. MicroPsi's motivational system reflects cognitive, social and physiological needs, and can account for individual variance and personality traits. Here, we describe the current state of the model that structures the behavior of cognitive agents in MicroPsi2.

Keywords: artificial general intelligence, MicroPsi2, motivation, motivational system, cognitive architectures

1 Introduction

MicroPsi [1] is an architecture for *Artificial General Intelligence*, based on a framework for creating and simulating cognitive agents [2]. Work on MicroPsi started in 2003. The current version of the framework, MicroPsi2 [3], is implemented in the Python programming language and may interface with various simulation worlds, such as Minecraft (see [4]). MicroPsi agents are hierarchical spreading activation networks that realize perceptual learning, motor control, memory formation and retrieval, decision making, planning and affective modulation.

One of MicroPsi's main areas of research concerns modeling a motivational system: whereas intelligence may be seen as problem solving in the pursuit of a given set of goals, human generality and flexibility stems largely from the ability to *identify* and *prioritize* suitable goals. An artificial system that is meant to model human cognition will have to account for this kind of autonomy. Our solution does not presuppose any goals, but instead a minimal orthogonal set of systemic *needs*, which are signaled to the cognitive system as *urges*. Goals are established as the result of learning how to satisfy those needs in a given environment, and to avoid their frustration. Since needs constantly change, the system will have to reevaluate its goals and behaviors continuously, which results in a dynamic equilibrium of activities. While cognition can change goals, expectations of reward and priorities, it cannot directly influence the needs itself.

MicroPsi's model of motivation [5] has its origins in the *Psi theory* [6, 7] and has recently been extended to account for a more detailed understanding of personality traits, aesthetic appreciation and romantic affection. While current MicroPsi2 agents do not implement all aspects of the motivational system (especially not the full set of need dimensions), the model is general enough to be adapted to a variety of applications, and has been integrated into other cognitive architectures, such as OpenCog [8]. In the following, I will focus especially on the discussion of concepts that can be transferred into other systems.

2 From Needs to Behavior

Since generally intelligent agents are not restricted to a fixed set of tasks and goals, they have to establish their own goals. These result from a pre-defined set of *needs*, or *demands*, which reflect what the agent has to consume, to achieve or to avoid, to thrive in its environment. Each need d reflects a variable c_d that varies between a target value v_d (the set-point at which the demand is fully satisfied) and a critical extreme v_0 (a point that may reflect where the system stops functioning). Whenever a need arises and no autonomous regulation is possible (for instance, metabolic regulation or transpiration to adjust the body temperature), the need is signaled by an *urge indicator.*

The difference between the target value and the current value is the *strength of the urge*: $urge_d = | v_d - c_d |$

The distance between the current value and the extremum corresponds to the *urgency* of satisfying the demand: $urgency_d = | v_d - c_d | \cdot | v_d - v_0 |^{-1}$

We could also specify that the range of a demand may extend below *and above* the set-point, so the urgency will have to reflect the distance from the current value to the nearest extremum. In practice, however, human demand regulation often uses two different urges for controlling upper and lower deviations from the set-point. For instance, heating and cooling, feeding and satiation, resting and exercising use different signals and control mechanisms.

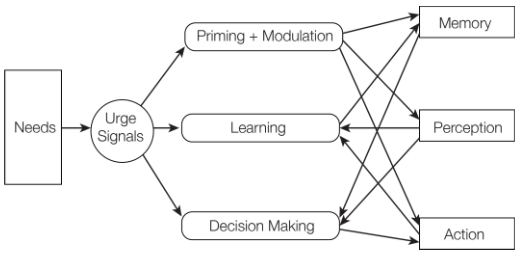

Fig. 1: Needs govern cognitive modulation, priming, learning and decision-making

Urge signals fulfill several roles in the cognitive system (figure 1). Their intensity governs arousal, execution speed and the resolution of cognitive processing (modulation). Changes in the urges indicate the satisfaction or frustration of the corresponding needs. Rapid changes are usually the result of an action of the agent, or of an event has happened to the agent. These changes are indicated with *pleasure signals* (corresponding to satisfaction) or *displeasure signals* (frustration). Pleasure and displeasure signals are used as reinforcements for *motivational learning* (figure 2). Each type of signal is connected to an *associator*. Associators establish a connection between two representations in MicroPsi; here, between the urge signal and the current situation or action. Furthermore, learning strengthens associations between the current situation and those that preceded it. In combination with mechanisms for forgetting and memory consolidation, this results in learning behavior sequences that end in goal situations. *Goals* are actions or situations that allow satisfying the urge (*appetitive goals*), or that threaten to increase the urge (*aversive goals*).

A *motive* is an urge that has been associated with a particular goal. Each action of a MicroPsi agent is based on an active motive, i.e. directed on reaching an appetitive goal, or avoiding an aversive goal.

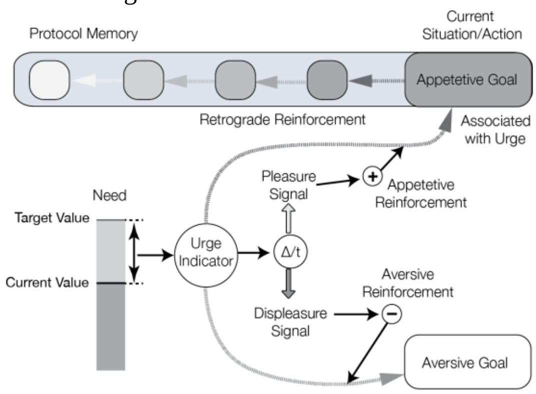

Fig. 2: Motivational learning

The association between urges and goals allows the agent to identify possible remedies whenever the need arises in the future, and *prime* its perception and memory retrieval towards the currently active urges. Most importantly, urges signal the current demands of the system, and thereby inform decision-making.

3 Types of Needs

The needs of a cognitive system fall into three groups: physiological needs, social needs and cognitive. Note that needs do not form a hierarchy, as for instance suggested by Maslow [9], but all act on the same level. This means that needs do not have to be satisfied in succession (many needs may never be fully satisfied), but concurrently. To establish priorities between different urges, each one is multiplied with a weight parameter $weight_d$ that expresses its importance relative to other needs. Each need has a $decay_d$ that specifies how quickly the c_d drops towards v_0 when left alone, and thus, how often it has to be replenished. Furthermore, each need has a parameter $gain_d$ and $loss_d$, which expresses how much it reacts to satisfaction or frustration. Using a suitable parameter set [*weight*, *decay*, *gain*, *loss*] for each demand or an agent, we can account for *individual variance of motivational traits*, and personality properties [10].

Physiological needs regulate the basic survival of the organism and reflect demands of the metabolism and physiological well-being. The corresponding urges originate in proprioceptive measurements, levels of hormones and nutrients in the bloodstream, etc. Physiological needs include sustenance (food and drink), physical integrity and pain avoidance, rest, avoidance of hypothermia and hyperthermia, and many more. (See figure 3 for an example implementation of physiological needs in MicroPsi2.)

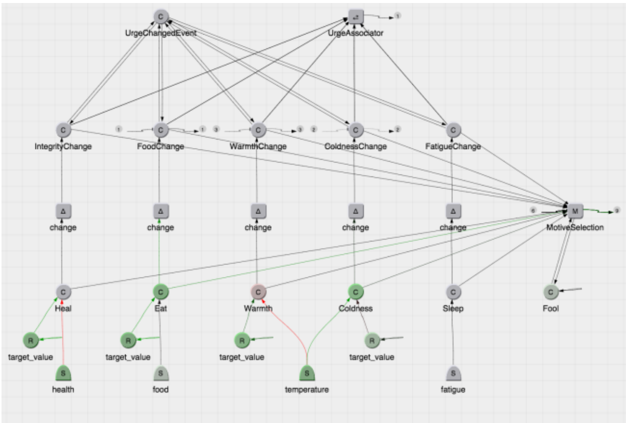

Fig. 3: Implementation of physiological needs in MicroPsi2 *(D. Welland)*

Social needs direct the behavior towards other individuals and groups. They are satisfied and frustrated by *social signals* and corresponding mental representations, but pleasure and displeasure from these sources is not necessarily less relevant to a subject than physiological pain. Consequently, people are often willing to sacrifice their food, rest, health or even their life to satisfy a social goal (getting recognition, supporting a friend, saving a child, winning a partner, maintaining one's integrity, avoiding loss of reputation etc.). Individual differences in the weight of social needs may result in more altruist or egotist, extraverted or introverted, abrasive or agreeable, romantic or a-romantic personalities.

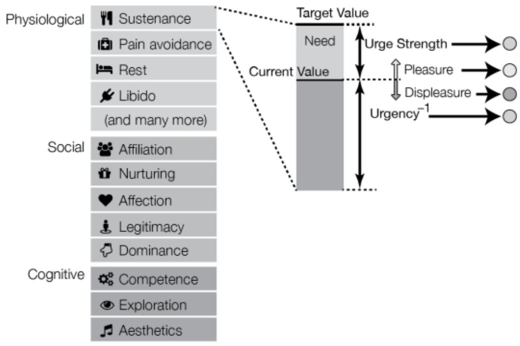

Fig. 4: Needs and motivator structure

Cognitive needs give rise to open-ended problem solving, skill-acquisition, exploration, play and creativity. Differences in the weight of cognitive needs may influence conscientiousness vs. spontaneity, openness, and hedonism.

Social Needs

Affiliation is the need for recognition and acceptance by other individuals or groups. It is satisfied by *legitimacy signals*, such as smiles and praise, and frustrated by frowns and reproach (*anti-legitimacy signals*). The effect of a legitimacy signal depends on its perceived strength, and the social weight (reputation) that the agent attaches to its source. For instance, praise from a highly respected individual typically results in a larger gain than anonymous praise. Affiliation plays the role of a virtual currency, offering rewards for cooperative behavior, and punishment for defection. A high decay of affiliation results in a personality that requires frequent attention from the environment, and introversion might be the outcome of a low decay or weight of affiliation.

Internal Legitimacy (“honor”) is a variant of affiliation that is directed on the conformance to internalized social norms. It directs social behavior in the absence of direct observers. Here, the source of legitimacy signals is the agent itself.

Nurturing is the need to altruistically care for other individuals or groups. It gets satisfied by subjectively increasing the well-being of someone else. The amount of satisfaction derived from an act nurturing depends on the social weight that the agent attaches to the object of his efforts. Supporting a cherished person or group will result in greater satisfaction than giving to an anonymous recipient. Psychopathy may be the result of the absence of a need for nurturing. Repeated exchange of affiliation and nurturing with another individual results in bonding, i.e. motivational learning will strongly establish that individual as a target for the respective urges. *Philia*/friendship is expressed as the reciprocal and mutual exchange of affiliation and nurturing between two individuals – it does not require an additional mechanism.

Romantic affection is the need to form a bond with a specific individual, leads to courtship behavior, and signals inclination for parental investment. Unlike libido, which is the physiological need for sexual gratification with a range of partners, romantic affection is directed on closeness and exclusive identification with a single partner [11]. A high weight of romantic affection may result in a propensity for *limerence* (an intense and sometimes dysfunctional infatuation), while a low one generates an a-romantic personality. Love, like friendship, is not an individual need (nor is it a single emotion). Instead, it may be best understood as a set of states, which differ based on the elementary mechanisms that are involved or absent: affiliation, nurturing, romantic affection and libido.

Dominance is the need to rise within the social hierarchies, or to maintain one’s place. A high weight of dominance leads to competitive behavior, and may result in a greater willingness to take risks or exert aggression to ascend.

Cognitive Needs

Competence is either task-related, effect-related or general:

- *Epistemic*, or *task-related competence* measures success at individual skills. The execution of a skill, and acquisition of new skills lead to satisfaction; failure, or anticipation of failure to frustration.
- *Effect-related competence* measures the ability to exert changes in the environment, with a gain that is proportional to the size of the observed effect.

- *General competence* is a compounded measure of the ability to satisfy needs (including the acquisition of epistemic competence). The strength of the urge is used as a heuristic to reflect on general performance, and to predict success at unknown tasks. Low general competence amounts to a lack of confidence.

Anticipated or factual failure at tasks may lead to an urge to replenish competence, which may be achieved by picking another task with a high probability of success. This dynamic may contribute to procrastination.

Exploration is the need to acquire certainty about objects and processes in the environment. The confirmation of uncertain expectations of perception and action increases certainty (and thus reduces the need for exploration); violations of expectations or a lack of available predictions increase the need for certainty. The strength of the urge for exploration is used as a heuristic for the degree of uncertainty in the given situation.

Stimulus-oriented aesthetics is the need for intrinsically pleasant stimuli, such as harmonious sounds, certain environmental features, tactile sensations. The favorable response to these stimuli is either a by-product of sensory processing, or serves indirect purposes (such as seeking out fertile land, identifying healthy mating prospects etc.).

Abstract aesthetics is the need to identify structure in mental representations, and to replace existing mental representations with more efficient ones. Abstract aesthetics is responsible for the discovery of mathematical elegance, musical arrangement, and for pleasure generated through the revision of conceptual structures.

4 Decision-Making

According to the Psi theory, all behaviors are either directed on the satisfaction of a need, or on the avoidance of the frustration of a need. Even serendipitous behavior is directed on need satisfaction (on exploration, rest or aesthetics). Identifying goals and suitable actions is the task of the decision-making system (figure 5).

Once a need becomes active and cannot be resolved by autonomous regulation, an urge signal is triggered, which brings the need to the attention of the *reactive* layer of the cognitive agent. Through past experiences, the urge has been associated with various actions, objects and situations that satisfied it in the past (appetitive goals), and situations and objects that frustrated it (aversive goals). Via activation spreading along these associations, relevant content in memory and perception is highlighted.

If an appetitive goal is perceived immediately, and there is no significant interference with current activity, the urge can be satisfied opportunistically, which will not require significant attentional processing. Otherwise, the agent attempts to suppress the new urge (by subtracting a selection threshold that varies based on the strength and urgency of the already dominant motive).

If the urge overcomes the selection threshold, and turns out to be stronger than the currently dominant urge, the current behavior is interrupted. The agent now tries to recall an applicable strategy, using spreading activation to identify a possible sequence of actions/world states from the current world situation (i.e., the currently active world model) to one of the highlighted appetitive goals. If such a strategy cannot be discovered automatically, the agent engages additional attentional resources and attempts to construct a plan, by matching known world situations and actions into a possible chain that can connect the current world situation to one of the appetitive goals. (At the moment, MicroPsi2 uses a simple hill-climbing planner, but many planning strategies can be used.)

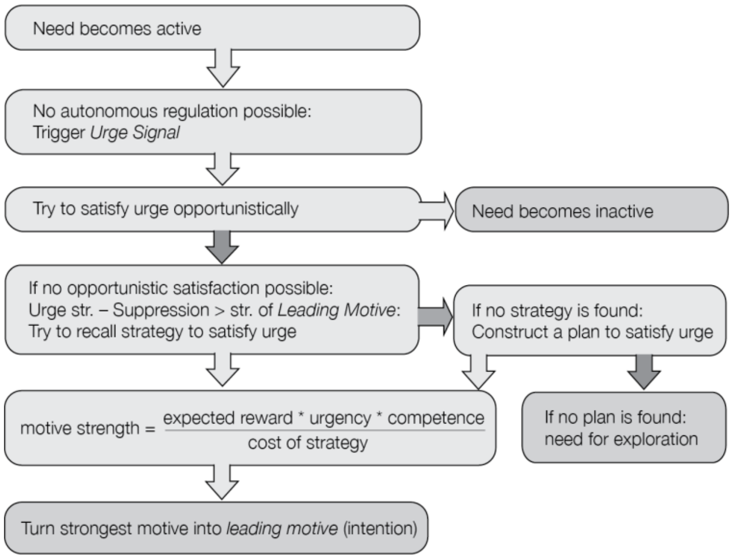

Fig. 5: Decision-making sequence

If plan construction fails, the agent gives up on pursuing the current urge, but increases its need for exploration (which will increase the likelihood of orientation behaviors to acquire additional information about the current situation, or even trigger experimental and explorative behavior strategies).

A successfully identified plan or automatism amounts to a motive (a combination of an active urge, a goal, and a sequence of actions to reach that goal). The strength of the motive is determined by estimating the reward of reaching the goal, the urgency of resolving the need, the probability of success, and dividing the result of these factors by the estimated cost of implementing the plan. The strongest motive will be raised to an intention, that is, it becomes the new dominant motive, and governs the actions of the agent. The probability of succeeding in implementing a strategy is currently estimated as the sum of the task specific competence (i.e. how likely the strategy succeeded in the past), and the general competence (to account for the agents general ability or inability to improvise and succeed in unknown circumstances).

5 Modulation

The strength of the needs of the agents does not only establish which goals an agent follows, but also *how* it pursues them. Cognitive and perceptual processing are configured by a set of global modulators (figure 6):

Arousal reflects the combined strength and urgency of the needs of the agent. Arousal is reflected in more energy expenditure in actions, action readiness, stronger responses to sensory stimuli, and faster reactions [12].

Valence represents a qualitative evaluation of the current situation. Valence is determined by adding all current pleasure signals to a baseline, and then subtracting all displeasure and currently active urges.

Aggression/submission determines the stance towards an attended object. Aggression is especially triggered by a negatively valenced reaction to another agent manifestly blocking a relevant goal (i.e., anger), and increases the likelihood of sanctioning behavior (fight). A low competence also leads to a low value of aggression, and reduces the inclination to engage (flight). In a more general sense, aggression suggests whether to approach or retract from the attended object, a middle value marks indifference.

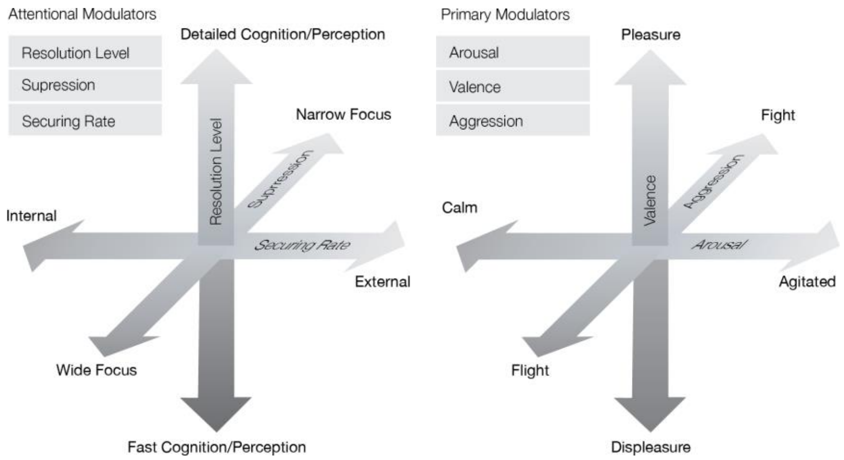

Fig. 6: Primary and attentional modulators

The combination of valence, arousal and aggression defines an *affective state* of the agent, and has originally been described by Wundt [13], who called the third dimension *tension*, and argued that each emotion can be characterized by their pleasurableness, intensity and stressfulness. The Wundt model of affect has been reinvented and modified numerous times in the history of psychology; the third dimension has also been identified as acceptance/rejection by Schlosberg [14] and submission/dominance by Mehrabian [15]. Note that arousal, valence and aggression are not themselves affects or emotions, but dimensions of a space of *cognitive configurations*. Affects are regions within that space. (Mehrabian calls this model *PAD*, for *pleasure*, *arousal* and *dominance*.)

MicroPsi uses currently six cognitive modulators. In addition to the three dimensions discussed above, these are:

Resolution level, the level of detail when performing cognitive and perceptual tasks. A high resolution will consider more details and thus often arrive at more accurate solutions and representations, while a low resolution allows faster responses. In MicroPsi, the resolution level is interpreted as the *width of activation spreading* in neuro-symbolic representations.

Suppression has already been mentioned in the context of decision-making. This modulator defines a selection threshold, which amounts to a stronger focus on the current task, and a narrower direction of attention. Suppression is a mechanism to avoid oscillations between competing motives.

Securing rate determines the frequency of obtaining/updating information from the environment. A dynamic environment requires more cognitive resources for perceptual processing, while a static environment frees resources for deliberation and reflection. In other words, the securing rate determines the direction of attention: outwards, into the environment, or inwards, onto the mental stage.

The three additional modulator dimensions configure the attention of a MicroPsi agent, by determining its width/detail, its focus, and its direction.

The values of the modulators are determined by the configurations of the urges, and by interaction among the modulators themselves (figure 7). Arousal is determined by the strength and urgency of all needs. A high arousal will also increase the resolution level and increase the suppression. The resolution level is increased by the strength of the current motive, but reduced by its urgency, allowing for faster responses. Suppression is increased by the strength and urgency of the currently leading motive, and is reduced by a low general competence. The securing rate is decreased by the strength and urgency of the leading motive, but increases with low competence and a high need for exploration (which is equivalent to experienced uncertainty). Aggression is triggered by agents or obstacles that prevent the realization of an important motive, and reduced by low competence.

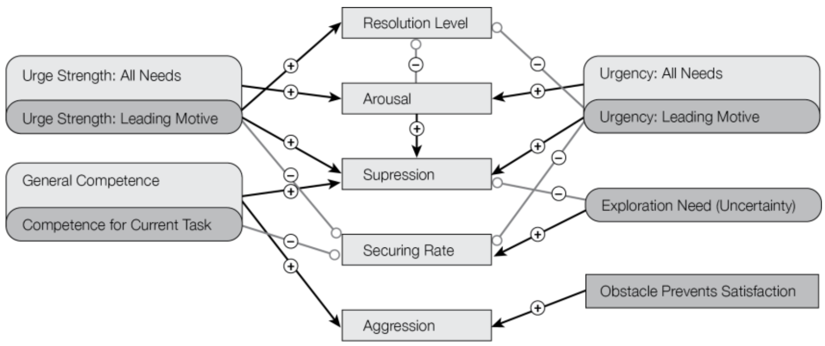

Fig. 7: The dynamics of modulation

Additionally, each modulator has at least four more or less fixed parameters that account for individual variance between subjects: the *baseline* is the default value of the modulator; the *range* describes the upper and lower bound of its changes, the *volatility* defines the reaction to change, and the *duration* describes the amount of time until the modulator returns to its baseline.

Emotions are either undirected, and can be described as typical configurations of the modulators, along with competence and experienced uncertainty, or they are a valenced reaction to an object, i.e. a particular motivationally relevant mental representation, combined with an affective state. Examples of undirected emotions are joy (positive valence and high arousal), bliss (positive valence and low arousal) or angst (negative valence, high experienced uncertainty, submission and low competence). Directed emotions are fear (negative valence directed on an aversive goal, submissiveness and low competence) and anger (negative valence directed on an agent that prevented an appetitive goal or caused the manifestation of an aversive goal, aggression, high arousal). Jealousy may either manifest as a fear (directed on losing romantic attachment or affiliation; submission), or as aggression (directed on an agent that prevents satisfaction of affiliative or romantic needs).

6 Summary

MicroPsi explores the combination of a neuro-symbolic cognitive architecture with a model of autonomous, polytelic motivation. Motives result from the association of urges with learned goals, and plans to achieve them. Urges reflect various physiological, social and cognitive needs. Cognitive processes are modulated in response to the strength and urgency of the needs, which gives rise to affective states, and allows for the emergence of emotions.

The current incarnation, MicroPsi2, adds further details to this motivational model, especially a more detailed set of social needs (nurturing, dominance and romantic affection). Parameters for each need (weight, gain, loss and decay) account for individual variation and modeling of personality traits. Modulators reflect valence, arousal and fight/flight tendency, as well as detail, focus and direction of attention. Modulators are parameterized by baseline, range, volatility and duration. We are currently applying the MicroPsi motivation model for analyzing the behavior of human subjects in computer games. The motivation model is also used to control behavior learning of autonomous AI agents in simulated environments.

While MicroPsi agents are implemented as hierarchical spreading activation networks, the underlying theory of motivation can be integrated into other cognitive models as well.

Acknowledgements: The implementation of MicroPsi would not be possible without the contributions of Ronnie Vuine, Dominik Welland and Priska Herger. I am grateful for generous support by and discussions with Dietrich Dörner, Martin Nowak and Jeffrey Epstein. Current work on MicroPsi is supported by the Program of Evolutionary Dynamics at Harvard University, and the Playful Systems Group at the MIT Media Lab.

References

1. Bach, J.: Principles of Synthetic Intelligence – An architecture of motivated cognition. Oxford University Press (2009)
2. Bach, J., Vuine, R.: Designing Agents with MicroPsi Node Nets. In Proceedings of KI 2003, Annual German Conference on AI. LNAI 2821, Springer, Berlin, Heidelberg (2003) 164-178
3. Bach, J.: MicroPsi 2: The Next Generation of the MicroPsi Framework. Proceedings of the Fifth Conference on Artificial General Intelligence (AGI 2012), Oxford, UK (2012) 11-20
4. Short, D.: Teaching Scientific Concepts Using a Virtual World—Minecraft. Teaching Science, 58 (3) (2012) 55-58
5. Bach, J.: A Framework for Emergent Emotions, Based on Motivation and Cognitive Modulators. International Journal of Synthetic Emotions (IJSE), 3(1) (2012) 43-63
6. Dörner, D.: *Bauplan für eine Seele*. Reinbeck (1999)
7. Dörner, D., Bartl, C., Detje, F., Gerdes, J., Halcour, D.: Die Mechanik des Seelenwagens. Handlungsregulation. Verlag Hans Huber, Bern (2002)
8. Cai, Z., Goertzel, B., Zhou, C., Zhang, Y., Jiang, M., Yu, G.: OpenPsi: Dynamics of a computational affective model inspired by Dörner's PSI theory. Cognitive Systems Research 17-18 (2012) 63-80
9. Maslow, A., Frager, R., Fadiman, J.: Motivation and Personality. (3rd edition.) Boston: Addison-Wesley (1987)
10. Bach, J.: Functional Modeling of Personality Properties Based on Motivational Traits. Proceedings of ICCM-7, International Conference on Cognitive Modeling, Berlin, Germany (2012) 271-272
11. Fisher, H. E.: Lust, attraction and attachment in mammalian reproduction. Human Nature, 9(1) (1998) 23-52
12. Pfaff, D. W.: Brain Arousal and Information Theory: Neural and Genetic Mechanisms. Cambridge, MA: Harvard University Press (2006)
13. Wundt, W.: Gefühlselemente des Seelenlebens. In: Grundzüge der physiologischen Psychologie II. Leipzig: Engelmann (1910)
14. Schlosberg, H. S.: Three dimensions of emotion. Psychological Review 1954, 61 (1954) 81-8
15. Mehrabian, A.: Basic dimensions for a general psychological theory. Oelgeschlager, Gunn & Hain Publishers (1980) 39–53