



City Research Online

City St George's, University of London

Citation: Besold, T. R., Kuehnberger, K-U., Garcez, A., Saffiotti, A., Fischer, M. H. & Bundy, A. (2015). Anchoring Knowledge in Interaction: Towards a Harmonic Subsymbolic/Symbolic Framework and Architecture of Computational Cognition. Lecture Notes in Computer Science, 9205, pp. 35-45. doi: 10.1007/978-3-319-21365-1_4

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/18668/>

Link to published version: https://doi.org/10.1007/978-3-319-21365-1_4

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Anchoring Knowledge in Interaction: Towards a harmonic subsymbolic/symbolic framework and architecture of computational cognition

Tarek R. Besold¹, Kai-Uwe Kühnberger¹, Artur d’Avila Garcez², Alessandro Saffiotti³, Martin H. Fischer⁴, and Alan Bundy⁵

¹ Institute of Cognitive Science, University of Osnabrück, {tbesold|kkuehnbe}@uos.de

² City University London, a.garcez@city.ac.uk

³ Örebro University, asaffio@aass.oru.se

⁴ University of Potsdam, martin.f@uni-potsdam.de

⁵ University of Edinburgh, a.bundy@ed.ac.uk

Abstract. We outline a proposal for a research program leading to a new paradigm, architectural framework, and prototypical implementation, for the cognitively inspired anchoring of an agent’s learning, knowledge formation, and higher reasoning abilities in real-world interactions: Learning through interaction in real-time in a real environment triggers the incremental accumulation and repair of knowledge that leads to the formation of theories at a higher level of abstraction. The transformations at this higher level filter down and inform the learning process as part of a permanent cycle of learning through experience, higher-order deliberation, theory formation and revision.

The envisioned framework will provide a precise computational theory, algorithmic descriptions, and an implementation in cyber-physical systems, addressing the lifting of action patterns from the subsymbolic to the symbolic knowledge level, effective methods for theory formation, adaptation, and evolution, the anchoring of knowledge-level objects, real-world interactions and manipulations, and the realization and evaluation of such a system in different scenarios. The expected results can provide new foundations for future agent architectures, multi-agent systems, robotics, and cognitive systems, and can facilitate a deeper understanding of the development and interaction in human-technological settings.

1 A Harmonic Analogy

Natural agents in many situations in their reasoning seem to rely on an enormous richness of representations (multimodal, grounded, embodied and situated), with many layers of representation at different levels of abstraction, together with dynamic reorganization of knowledge. Also, real-world situations require agents to perform what can be interpreted as dynamic changes or alignments of representation, as different agents might use different languages and levels of description. Unfortunately, when trying to follow the natural example by transferring and (re)creating this representational richness and diversity to artificial agents, the resulting mismatches cannot be cured by standardization, but arise due to differences in the environment, tasks to be solved, levels of abstraction, etc. Additionally, real-world applications also demand online and

bidirectional learning that takes place in real-time, as well as the adaptation to changes in the environment, to the presence of new agents, and to task changes.

A conceptually similar situation presents itself in the domain of music: Music appears on different levels as there are, among others, a physical level (audio data), a MIDI level, a chord progression level, a harmonic, melodic, rhythmic level, a score level, a structural level of a piece of music, a (semantic) meta-level for describing music. Concerning the interaction and transfer of information between levels, in certain cases there are obvious mappings (e.g. MIDI level to score to harmonic structure), in others there are partial or incomplete mappings (e.g. harmonic structure to score, rhythmic to physical level), in others there are fuzzy or tentative mappings (e.g. melody to harmony (in an idiom) or to rhythmic level (in an idiom), physical to structural level of a piece of music), and between others there are no mappings at all (e.g. MIDI level to semantic/meta level, melodic level to structural to harmonic level). Also, music can be described in different representation formats on all levels. A piece of music can then be considered as a multi-layered multi-representational entity with certain connections and constraints (in form of relations, mappings etc.) between the layers, where, for instance, changing the chord progression influences (in an obvious, partial, or fuzzy way) many (but not all) other levels.

From a functional perspective, pieces of music that have been analyzed in such a multi-representational way could, among others, be used to learn or to detect obvious mappings between the layers, to detect novelties and correlations, to systematically unfold the specific properties of pieces (or classes thereof)/idioms/genres of music, or to find the invariant properties of music (e.g. a change of melody changes systematically the score, but does not affect the larger structure of the piece).

Returning to the agent setting by way of analogy we envision a system operating on different levels of representations (corresponding to different formal layers in the system's architecture) similar to the musical case. The hierarchy could consist, for instance, of a (lowest) neural layer learning on the perception/motor level, an anchoring layer learning elementary (semi-)symbolic representations of objects, a reactive layer taking over in critical situations, a deep learning layer learning on more abstract levels, a symbolic layer doing reasoning and planning, and a (higher) symbolic layer providing the core ontology. Like in music, some of these layers have obvious, some have partial, some have fuzzy, and some have no mappings/relations between themselves.

Now, a corresponding architecture should be in a "pre-established" harmony: Triggering an abstract plan to move from A to B should result in the motor action to move from A to B, classifying on the neural level a certain perceptual input such as, for instance, a chair should result in the activation of the concept "chair" in the ontology or the working memory, and so on. And whilst the basic links might be hard coded, learning a new concept on the subsymbolic level should somehow result in a new concept entry in the ontology, i.e., there should be interaction between the different layers in terms of information and conceptualizations. Finally, when thinking about a simulated or actual system that is operating on these interacting levels in a multi-representational manner it should allow for similar mechanisms and interactions as in the music case.

2 The Core Ideas

Addressing the challenges outlined in the previous section and taking inspiration in the sketched analogy to the musical domain, we propose the development of a new approach and integrated techniques that will enable the sustainable and accessible creation of large-scale integrated knowledge repositories for use by multi-agent systems or as part of a cyber-physical system. In this note, we suggest a research program for the community working on embedded intelligence. This program for 'anchoring knowledge in interaction', aims at developing, theoretically and practically, a conceptual framework and corresponding architecture that model an agent's knowledge, thinking, and acting truly as interrelated parts of a unified cognitive capacity. That is, knowledge is seen as multi-layered phenomenon that appears at different levels of abstraction, promotes interaction between these levels of abstraction, is influenced by the interaction between agent and environment (potentially including other agents), and is essentially linked to actions, perception, thinking, and being. The program's long term vision, thus, is a radically new paradigm in what concerns interaction styles (which are action-centered, embodied, multi-modal), knowledge repositories (with different levels and forms of knowledge representation, as, e.g., multi-modal, hybrid), and user modeling and communication through learning and adaptation.

The scientific aims of the described endeavor target advances at different conceptual and topical levels (covering, among others, all three levels of analysis of a cognitive system described in [19]). On the embodiment level, it shall be shown that elementary forms of representations can be learned from an agent's interactions within an environment. The resulting multi-modal representations may be noisy, they may be uncertain and vague, it may be the case that different agents have different languages for representing knowledge, or that changes in the environment may come into play. On this level, building on recent advances in the study of embodied cognition, the main development will therefore be an extension of the well-known anchoring framework in robotics [5] to grounding not only objects, but also certain general observable properties appearing in the environment.

The embodiment view of knowledge provides an interaction-based neural representation of knowledge that is not represented at the conceptual level. Neural systems can promote robust learning from data, as part of an online learning and reasoning cycle to be measured in terms of an improved experience, a faster adaptation to a new task, and the provision of clear descriptions. On this level, a lifting procedure shall be specified that will produce descriptions, thus lifting grounded situations and an agent's action patterns to a more abstract (symbolic) representation, using techniques from machine learning like deep networks and analogy-making. This can be seen as a natural consequence of recent research developed for deep learning and neural-symbolic computing, the crucial added value over the state of the art being the combination of these new methodologies with analogical transfer of information between representation systems.

Knowledge at a symbolic level is usually considered to be static and error-intolerant. Due to the fact that initial multi-modal representations lifted from the subsymbolic level can be error-prone, and that different agents might use different and a priori possibly incompatible representation languages, the program's objective at the level of symbolic representations is a dynamic re-organization based on ontology repair mechanisms,

analogy, concept invention, and knowledge transfer. These mechanisms foster adaptation of an agent to new situations, the alignment between representations of different agents, the reformulation of knowledge entries, and the generation of new knowledge.

In summary, the envisioned account of the emergence of representations through cognitive principles in an agent (or multi-agent) setting can be conceptualized as follows: Grounding knowledge in cognitively plausible multimodal interaction paradigms; lifting grounded situations into more abstract representations; reasoning by analogy and concept blending at more abstract levels; repair and re-organization of initial and generated abstract representations.

Applications for such a framework are manifold and not limited to the “classical” realm of robotic systems or other embodied artificial agents. Also, for instance, future tools in e-learning education – in order to guarantee sustainable and life-long learning tools for different groups of learners – will focus on aspects such as, for instance, adaptivity to target groups of learners, modeling of the knowledge level of group members, multi-modality, integration of a richer repertoire of interaction styles of learning including action-centered set-ups, promotion of cooperative and social learning, etc. Such devices are inconceivable without a cognitive basis, adaptation, multiple representations, concept invention, repair mechanisms, analogical transfer, different knowledge levels, and robust learning abilities.

3 The Core Objectives

The core idea is that knowledge is multi-layered, i.e. there is no static, fixed, and definite representation of knowledge, rather agents have to adapt, learn, and re-organize knowledge continuously on different levels while interacting with other agents and their environment. Thus, the future architecture aims to anchor and embody knowledge by the interaction between the agent and its environment (possibly including other agents), to give an approach to lift the resulting situated action patterns to a symbolic level, to reason by analogy on the abstract and the subsymbolic level, to adapt, or in case of clashes, repair the initial representations in order to fit to new situations, and to evaluate the approach in concrete settings providing feedback to the system in a reactive-adaptive evolutionary cycle.

The project’s scope is primarily focused on providing answers to several longstanding foundational questions. Arguably the most prominent among these, together with answers based on the conceptual commitments underlying the discussed research program, are:

- 1.) *How does knowledge develop from the concrete interaction sequences to the abstract representation level?* The crucial aspect is the lifting of grounded situations to more abstract representations.
- 2.) *How can experience be modeled?* Experience can be explained by deep learning.
- 3.) *How is deeper understanding of a complex concept made possible?* Theory repair makes precisely this possible.
- 4.) *To which extent do social aspects play a role?* Analogical transfer of knowledge between agents is a central aspect concerning efficient and flexible learning and understanding.

Although efforts are directed towards reintegrating the different aspects of agent cognition spanning from abstract knowledge to concrete action, there is also a strong drive toward new concepts and paradigms of cognitive and agent-based systems. A fresh look at the embodiment problem is proposed, as the envisioned account goes significantly beyond the perception-action loop and addresses the problem of the possibility of higher intelligence where it occurs, namely at the level of the emergence of abstract knowledge based on an agent's concrete interaction with the environment. Similarly, learning aspects are tackled not only on a technical level, but furthermore pushed beyond the technical area by gaining inspiration from cognitive science and concept-guided learning in the sense of analogical learning and concept blending, as well as from newer findings in neural networks learning.

4 Structure and Methods

The new approach for modeling knowledge in its breadth, namely from its embodied origins to higher level abstractions, from the concrete interaction between an agent and its environment to the abstract level of knowledge transfer between agents, and from the holistic view of knowledge as an interplay between perception, (inter)action, and reasoning to specific disembodied views of knowledge, touches on different aspects and fields of research. It therefore requires the integration of expressive symbolic knowledge representation formalisms, relational knowledge, variables, and first-order logic on the one hand with representations of sensorimotor experiences, action patterns, connectionist representations, and multi-modal representations on the other.

The different topics above will be formalized, algorithmically specified, implemented in running applications and evaluated. With respect to the formalization, research methods from machine learning (e.g. cross-validation [9] or layer-wise model selection [1] in deep networks) will be used to learn conceptual knowledge from sub-symbolic data, i.e. to extract knowledge from such networks in order to lift and enable transfer learning on the conceptual level. This type of conceptual knowledge will be used as input to the analogy-making process to generate new concepts by abstraction and transfer of knowledge in a domain-independent and multi-modal setting. The formalization of the analogy process, including the computation of generalizations [21], and the multi-modal embodied representations potentially change the signatures of the underlying language(s). Therefore, the theory of institutions [8] will be used as methodology in which dynamic changes of languages can be rigorously formalized. The repair of theories and concept invention mechanisms will be linked to analogy-making and are methodologically formalized in a higher-order logical framework [3, 17].

The corresponding research program is structured into interrelated thrusts:

1.) Cognitive Foundations of Knowledge: New embodied approaches to understanding human cognition, augmenting the traditional symbol manipulation-based accounts, emphasize the importance of sensorimotor interactions as part of knowledge formation [10]. Thereby, they provide the starting point for a systematic assessment of basic learning signatures in the presence of different sensorimotor experiences, leading to recommendations for the development of cognitively-inspired formal frameworks

for embodied computation, in particular, for the specification of learning mechanisms, analogy, and repair mechanisms.

Together with approaches from computational neuroscience and network-level cognitive modeling (as, e.g., the recently proposed framework of conceptors in dynamical system models [15]) work in this thrust will create the cognitively-inspired foundations and low-level input representations and content for the subsequent stages of processing and reasoning.

2.) Anchoring Knowledge in Perception, Action, and Interaction: Anchoring [5] in robotic systems is the problem of how to create, and to maintain in time and space the connection between the symbol- and the signal-level representations of the same physical object. Anchoring this far is concerned with the grounding of symbols that refer to specific object entities, i.e. anchoring can be considered as a special case of the symbol grounding problem limited to physical objects.

While different approaches to solving this foundational problem have been proposed [4], a satisfactory answer is still elusive and the arising difficulties are manifold: In a distributed system, individual agents may need to anchor objects from perceptual data coming either from sensors embedded directly on the robot or information coming from external devices. Further, agents each with their own anchoring module may need to reach a consensus in order to successfully perform a task in a cooperative way. Also Human-Robot Interaction (HRI)-oriented communication about objects requires a coordinated symbol-percept link between human and robot.

In the envisioned framework, building on [16]’s results on cooperative anchoring and on [6]’s symbiotic HRI robotic systems, anchoring happens under even more general conditions: anchoring is performed both top-down and bottom-up during learning; new symbols for new objects and categories are dynamically introduced by repair and concept invention mechanisms; the denotation of a symbol used in communication must be consistent across the communicating agents; anchoring must enable the establishment of analogical links across different agents.

3.) Lifting Knowledge from the Subsymbolic to the Symbolic Level: The elementary forms of representations referred to above, which may be noisy, vague, and uncertain, have been made suitable for learning through the use of neural networks, notably recently deep networks.

Deep learning is a form of representation learning aiming at discovering multiple levels of representation. Also, recent advances in the area of deep learning have shown promising results when applied to real-time processing of multimodal data [7], and state-of-the-art deep learning methods and algorithms have been able to train deep networks effectively when applied to different kinds of networks, knowledge fusion, and transfer learning [2]. However, more expressive descriptions and forms of representation have become more difficult to obtain from neural networks.

Following a neural-symbolic approach, neural learning will be combined with temporal knowledge representation using variations of the Restricted Boltzmann Machine model [14]. The resulting approach will offer a method for validating hypotheses through the symbolic description of the trained networks whilst robustly dealing with uncertainty and errors through a Bayesian inference model. Furthermore, the use of Gärdenfors’ “conceptual spaces” [11] to link symbolic and subsymbolic data, as done in [16],

will also be investigated and tested for its applicability and feasibility in the proposed complex sensing, learning, and reasoning cycle.

4.) Analogy/Blending: Analogy is classically understood as a method to detect and operate on structural commonalities between two domains, and in cognitive science and cognitive AI has been applied to a variety of tasks, e.g. intelligence tests [18], learning with sketches [20], or naive physics [21]. Unfortunately, until now analogy engines are designed only for highly specialized domains, neither multi-modal representations nor embodied interaction with the environment is taken into account, abstraction and knowledge projection from source to target are usually restricted to a few stages of analogical comparisons, and repair strategies for faulty inputs are rather limited.

The described approach brings analogical reasoning from a computer science perspective closer to its cognitive origins: generalizability, multi-modal representations, and embodied interaction with the environment are considered to be essential for analogy-making in this project. Furthermore, analogies will directly be linked to repair mechanisms in order to facilitate the resolution of errors. Thus, analogies are re-considered concerning their foundations and re-conceptualized concerning their methodological basis, as well as their applications.

5.) Concept Formation/Reformation: An important way in which new concepts are formed is by the evolution of existing concepts that have proved inadequate: Such inadequacies are often revealed by failures of inference using the old concepts. Researchers lately explored how these inadequacies can trigger conceptual change in different domains as, e.g., physics [17] or in ontologies [12].

The resulting domain-specific diagnosis and repair mechanisms bore strong similarities to each other: The so called reformation algorithm (a modification of unification) is an attempt to capture the generality behind these mechanisms and provide a domain-independent diagnosis and repair mechanism for conceptual change (cf. [3] for an example). Based on this approach, generic mechanisms for repairing agents' faulty representations (especially those produced by imperfect analogies) will be developed, implemented, and evaluated in a variety of domains going far beyond current (domain specific) solutions.

5 First Steps Towards an Implementation

At the current stage, the suggested research program is still mostly in its conception and planning phase. Nonetheless, a basic conceptual architecture (see Fig. 1) can already be laid out based on the considerations discussed in the previous sections: depending on the perspective and degree of abstraction, this architecture can either be sub-divided into five hierarchical layers (respectively corresponding to the five thrusts sketched in the previous section) or can be conceptualized as structured in three (partially overlapping) functional components. In the latter case, the cognitive foundations and the anchoring layer are combined into a low-level subsymbolic module, analogy/blending and concept formation/repair into a high-level symbolic module, and anchoring, knowledge lifting, and analogy into an intermediate module bridging in the direction from the low-level to the high-level component. Concerning the individual modules, interaction happens both between layers within components (as, e.g., between analogy/blending and con-




Fig. 1. An overview of the conceptual structure, functional components, and the interplay between layers of the envisioned architecture implementing the cycle of learning through experience, higher-order deliberation, theory formation and revision.

cept formation/reformation layer) as well as across components (as, e.g., through the feedback from the concept formation/reformation to the anchoring). This results in an architecture adhering to and implementing the “harmonic analogy” setting from the introductory section, with changes in one layer propagating to others in order to maintain a “harmonic” configuration.

Within the low-level module, conceptors and similar approaches are employed in order to establish a certain initial structure of the perceptual input stream on a subsymbolic level, additionally reinforcing the proto-structure already imposed by the properties of the embodiment-inspired approach to computation. This initial structure can then be used as basis upon which the anchoring layer operates, coupling elements of this structure to objects and entities in the perceived environment and/or to action-based percepts of the agent. This coupling goes beyond the classical accounts of anchoring in that not only correspondences on the object/entity level are created, but also properties and attributes of objects/entities are addressed. Thus, subsymbolic correspondences between the initial structured parts of the perceptual input stream as representational vehicles and their actual representational content are established. These vehicle-content pairs then can be arranged in a hierarchical structure, both on object/entity level and on connected object/entity-specific property levels, based on general attributes of the perceptual input stream (as, e.g., order of occurrence of the respective structures, relations between structures) hinting at elements of the representational content, and on direct properties of the representations in their function and form as representational vehicles.

Within the high-level module, analogy and blending are applied on rich logic-based representations to find corresponding concepts and knowledge items, to transfer and adapt knowledge from one context into an analogically-related similar one, and to combine existing concepts into new concepts based on analogical correspondences between the inputs. Still, these processes are error-prone in that they can reveal inconsistencies between existing concepts, or can introduce new inconsistencies by concept combination or concept transfer and adaptation. Arising inconsistencies can then be addressed by the top-level concept formation and reformation layer, allowing to repair inconsistent symbolic representations through manipulations of the representational structure and to introduce new representations or concepts by introducing new representational elements – and, when doing so, informing and influencing the subsymbolic anchoring layer to perform corresponding adaptations in its vehicle-content correspondences.

Finally, the intermediate module bridging from low-level to high-level processing takes the correspondences between representing structures and representational content established by the anchoring layer, and uses deep learning techniques for representation learning in order to lift the subsymbolic vehicle-content pairs to a logic-based form of representation. Here, the corresponding learning process will take into account already existing knowledge on the symbolic side by way of analogy both, over vehicle-content pairs and over the learning process itself (i.e., resulting in a form of cross-informed transfer learning): When presuming a (fairly low) basic level of continuity of the environment and the perceptual input stream, on the one hand, over time the symbolic forms of newly lifted vehicle-content pairs most likely will share analogical commonalities with already existing concept and knowledge items which can then be used to foster the lifting process, while on the other hand successive or parallel lifting processes also can cross-inform each other leveraging the analogical structure over processes and exploiting shared or similar sub-parts.

6 (Far) Beyond Multi-Level Data Fusion

At first sight, similarities between the proposed research project and work in multi-level data fusion might be suggested, questioning the sketched approach's novelty or added value over existing accounts.

Still, the differences are significant. Data fusion tries to leverage the advantage of receiving several data streams concerning the same source for getting a more precise characterization of the source: *“data fusion techniques combine data from multiple sensors and related information from associated databases to achieve improved accuracy and more specific inferences than could be achieved by the use of a single sensor alone.”* [13]. Even when leaving aside the targeted improvements and extensions to existing techniques, such as performing anchoring also on the attribute level, the ambition of the research project sketched in this paper goes far beyond this: The final goal is the development of a cognitively-inspired combination of low-level sensing with high-level reasoning in an attempt of anchoring (symbolic) knowledge in (subsymbolic) perception and (inter)action in a continuous feedback loop.

If successful, this would in all likelihood constitute a significant step towards the (re)creation of the foundation for cognitive capacities and forms of reasoning in next

generation systems in artificial intelligence, as well as major progress towards developing a computational test bench and agent model for theories from cognitive science.

References

1. Arnold, L., Paugam-Moisy, H., Sebag, M.: Unsupervised Layer-Wise Model Selection in Deep Neural Networks. In: *Proceedings of ECAI 2010: 19th European Conference on Artificial Intelligence*. pp. 915–920. IOS Press (2010)
2. Bengio, Y.: Learning Deep Architectures for AI. *Foundations and Trends in Machine Learning* 2(1), 1–127 (2009)
3. Bundy, A.: The interaction of representation and reasoning. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 469(2157) (2013)
4. Chella, A., Frixione, M., Gaglio, S.: Anchoring symbols to conceptual spaces: the case of dynamic scenarios. *Robotics and Autonomous Systems* 43(2–3), 175–188 (2003)
5. Coradeschi, S., Saffiotti, A.: Anchoring symbols to sensor data: Preliminary report. In: *Proceedings of the 17th AAAI Conference*. pp. 129–135. AAAI Press (2000)
6. Coradeschi, S., Saffiotti, A.: Symbiotic Robotic Systems: Humans, Robots, and Smart Environments. *IEEE Intelligent Systems* 21(3), 82–84 (2006)
7. De Penning, H.L.H., Garcez, A.S.D., Lamb, L.C., Meyer, J.J.C.: A Neural-symbolic Cognitive Agent for Online Learning and Reasoning. In: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*. pp. 1653–1658. AAAI Press (2011)
8. Diaconescu, R.: *Institution-independent Model Theory*. Birkhäuser, 1st edn. (2008)
9. Dietterich, T.G.: Approximate Statistical Tests for Comparing Supervised Classification Learning Algorithms. *Neural Comput.* 10(7), 1895–1923 (1998)
10. Fischer, M.H.: A hierarchical view of grounded, embodied, and situated numerical cognition. *Cognitive Processing* 13(1), 161–164 (2012)
11. Gärdenfors, P.: *Conceptual Spaces: The Geometry of Thought*. MIT Press (2000)
12. Gkaniatsou, A., Bundy, A., Mcneill, F.: Towards the Automatic Detection and Correction of Errors in Automatically Constructed Ontologies. In: *8th International Conference on Signal Image Technology and Internet Based Systems 2012*. pp. 860–867 (2012)
13. Hall, D., Llinas, J.: An introduction to multisensor data fusion. *Proceedings of the IEEE* 85(1), 6–23 (1997)
14. Hinton, G.E.: A Practical Guide to Training Restricted Boltzmann Machines. In: Montavon, G., Orr, G., Müller, K.R. (eds.) *Neural Networks: Tricks of the Trade, Lecture Notes in Computer Science*, vol. 7700, pp. 599–619. Springer (2012)
15. Jaeger, H.: Controlling recurrent neural networks by conceptors. *arXiv* (2014), 1403.3369v1 [cs.CV], 13 Mar 2014
16. LeBlanc, K., Saffiotti, A.: Cooperative anchoring in heterogeneous multi-robot systems. In: *2008 IEEE International Conference on Robotics and Automation*. pp. 3308–3314 (2008)
17. Lehmann, J., Chan, M., Bundy, A.: A Higher-Order Approach to Ontology Evolution in Physics. *Journal on Data Semantics* 2(4), 163–187 (2013)
18. Lovett, A., Forbus, K., Usher, J.: A structure-mapping model of Raven’s Progressive Matrices. In: *32nd Annual Meeting of the Cognitive Science Society*. pp. 2761–2766 (2010)
19. Marr, D.: *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman and Company (1982)
20. McLure, M., Friedman, S., Forbus, K.: Learning concepts from sketches via analogical generalization and near-misses. In: *32nd Annual Meeting of the Cognitive Science Society*. pp. 1726–1731 (2010)
21. Schwering, A., Krumnack, U., Kühnberger, K.U., Gust, H.: Syntactic Principles of Heuristic-Driven Theory Projection. *Journal of Cognitive Systems Research* 10(3), 251–269 (2009)