Equipping Symbolic Frameworks with Soft Computing Features

Kai-Uwe Kühnberger
Institute of Cognitive Science, University of Osnabrück
kkuehnbe@uos.de

Abstract

This paper proposes to have a fresh look on the neural-symbolic distinction by focusing on the strengths and weaknesses of the two antagonistic approaches. We claim that in both worlds, the symbolic and the subsymbolic world, there is a tendency to embrace new methods borrowed from the respective other methodology. Whereas, this seems to be quite obvious from the neural perspective we focus on sketching ideas where soft computing methods are used in classical symbolic, logic-based frameworks. We exemplify this claim by some remarks concerning certain soft computing features of Heuristic-Driven Theory Projection (HDTP), a symbolic framework for analogy-making and concept blending.

1 Introduction

From a certain practical perspective the long lasting distinction between symbolic and subsymbolic approaches is rather natural. Strengths and weaknesses of these approaches are quite complementarily distributed as shown in Table 1. For example, whereas symbolic approaches have their strengths mostly in higher cognitive abilities, such as planning, reasoning, and the ability to represent knowledge explicitly, subsymbolic approaches are very successful in areas related to lower cognitive abilities, like data-driven approaches for modeling perception, learning, and adaptation to a rapidly changing environment. Therefore, it is a natural idea to develop frameworks that combine the (complementary) strengths of both approaches such that at the same time the (complementary) weaknesses of the two types of approaches are avoided.

Although, many frameworks for reconciling these two major types of models have been proposed,[1] the last decades of research in this direction has shown that the development of such integrated frameworks that show their strengths in a broad variety of classical applications is everything else than easy to achieve. To phrase this more frankly: Taking the last two decades of neural-symbolic integration into account, it is rather obvious that neither the ultimate application nor the ultimate theoretical insight has been proposed. Various frameworks have been studied theoretically and practically, strengths and weaknesses of the different accounts have been evaluated. Nevertheless, there is no agreement whether the proposed frameworks can do better than alternative (classical) approaches that are not based on the spirit of neural-symbolic reasoning and learning. Even worse many researchers would even say that neural-symbolic reasoning and learning devices have failed to demonstrate their broad applicability in theory and practice.

This overall negative claim is surprising because there are many proposals on the market. Just to mention some of these approaches towards neural-symbolic integration, the frameworks described in [Hitzler *et al.*, 2004], [Garcez *et al.*, 2002], and [Gust *et al.*, 2007] try to integrate complex higher reasoning abilities into a neural network inspired approach. Although, there are theoretically highly interesting results available, practical applications are rather rare. Besides the mentioned class of models, there is furthermore a large number of hybrid models in the field of cognitive architectures that work in similar directions towards a neural understanding of higher cognitive abilities with neurally inspired means. A good overview of hybrid systems can be found for example in [Wermter and Sun, 2000]. Unfortunately, hybrid cognitive architectures have often similar acceptance problems in parts of the scientific community as cognitive architectures in general are often confronted with: many classical researchers do not take them seriously into account, because specialized engineering solutions are just more successful in most practical applications.

We think that the situation in neural-symbolic reasoning and learning is unsatisfactory because of the described lack of theoretical and practical breakthroughs. It may be a good strategy to take a more abstract perspective on the interplay between subsymbolic and symbolic frameworks into account. In particular, it may be reasonable to focus on current tendencies of symbolic approaches to model learning and adaptation aspects one the one side, and the tendencies of subsymbolic approaches to represent complex data structures and higher cognitive processes on the other. This more abstract level of neural-symbolic integration, namely on the level of

[1] The collection [Hammer and Hitzler, 2007] gives a good overview concerning different frameworks and different methodologies in order to bridge the gap between subsymbolic and symbolic approaches.

	Symbolic Approaches	Sub–Symbolic Approaches
Methods	(Mostly) logical and/or algebraic	(Mostly) analytic
Strengths	Productivity, Recursion Principle, Compositionality	Robustness, Learning Ability, Parsimony, Adaptivity
Weaknesses	Consistency Constraints, Lower Cognitive Abilities	Opaqueness Higher Cognitive Abilities
Applications	Reasoning, Problem Solving, Planning etc.	Learning, Motor Control, Vision etc.
CogSci Relation	Not Biologically Inspired	Biologically Inspired
Other Features	Crisp	Fuzzy

Table 1: Strengths and weaknesses of symbolic and neural/subsymbolic approaches. The distributions of these strengths and weaknesses are quite complementary.

strengths and weaknesses of the respective accounts, comes together with an integration of new methods that are not considered as standard techniques in the symbolic and subsymbolic worlds. Nevertheless, the big difference to general approaches to bridge the gap between symbolic and subsymbolic models is the local character of the integration. Adding a certain feature of the subsymbolic world into symbolic frameworks (or vice versa) is not solving the general problem of the symbolic-subsymbolic distinction. The hope is that a large number of such integrations facilitates a general theory of neural-symbolic reasoning and learning.

This paper attempts to shed light on the underlying problem from a rather general perspective. First, we will describe (from a subjective perspective) current tendencies that symbolic and subsymbolic frameworks converge against each other (Section 2). By using the term relatively loosely, this convergence can be subsumed under the endeavor to develop neural-symbolic integration devices. In Section 3, we will describe informally certain soft computing features of the symbolic framework Heuristic-Driven Theory Projection (HDTP), i.e. we give an example which shows that certain symbolic approaches integrate features of the subsymbolic world quite obviously into their internal procedures. Last but not least, Section 4 adds some speculations of how the totality of these endeavors could eventually bring fresh ideas into the neural-symbolic integration research field and concludes the paper.

2 Convergence Tendencies of the Neural and Symbolic Worlds

We claim that there are certain convergence tendencies between symbolic and subsymbolic frameworks that attempt to model strengths of the respective alternative approach. In other words, there is the tendency of developing neural models that expand their applicability to higher cognitive abilities. On the other hand, there are also tendencies in the development of symbolic frameworks to model strengths of neurally inspired or soft computing approaches like learning abilities, adaptivity, and robustness with respect to noisy data. Whereas, for the neural approaches this seems to be quite obvious – e.g. the proceedings of the International Workshops for Neural-Symbolic Reasoning and Learning are a good resource for these developments[2] – this is not as clear for models that are based primarily on symbolic methodologies. Therefore, we will list informally some of these developments where symbolic approaches are intended to expand to the neural world.

- **Relational Learning:** The equipment of the framework of inductive logic programming with probabilistic features can be seen as an example of extending a classical logical (and therefore symbolic) learning approach with subsymbolic (or soft computing) methods. A standard reference for relational learning can be found in [De Raedt, 2008].
- **Ontology Repair System:** The computational modeling of scientific discovery requires relatively expressive and dynamically changing formalisms for computational approaches. An algorithmic approach for finding new insights in science, for inventing new scientific concepts, for re-interpreting old concepts in scientific theories newly and the like require a flexible formalism that allows not only extensions of a theory, but moreover changes of the underlying languages. A guiding idea of these approaches is to resolve clashes in existing scientific theories by changing domain theories and their languages in a non-trivial way. Compare [Bundy and Chan, 2008] and [Chan *et al.*, 2010] for more information concerning automated approaches towards ontology evolution in physics.
- **Dynamics of Analogy Making:** It is quite uncontroversial that a fundamental mechanism for cognition is analogy making [Hofstadter and the Fluid Analogies Research Group, 1995]. In particular, if cross-domain analogies are considered, then computational models for analogical reasoning require dynamic changes of the signatures of languages (provided we restrict our attention to symbolic approaches). Furthermore, such models compute ranked candidates for analogical relations, i.e. they add a classical soft computing feature of more or less plausible candidates for an analogy to the framework.
- **Ontologies in Text Understanding Systems:** In the context of language understanding systems there is the need to integrate not only linguistic knowledge into the systems's knowledge base, but also world knowledge, often represented in form of domain ontologies. Usually these ontologies are considered to be crisp. Nevertheless, newer approaches towards natural language understanding systems attempt to integrate soft computing elements into the symbolic representations of the integrated ontologies. A good examples for such a general multi-purpose system can be found in [Ovchinnikova, 2012].

These are just a few examples among many possible candidates showing that there are attempts to equip symbolic systems with soft computing features that allow adaptation, dynamic changes, conflict resolution, learning, and the invention of new concepts. Although most of these systems do not

[2] Cf. http://www.neural-symbolic.org/ for further reference.

implement a classical neural approach (in terms of a network of neurons, activation potentials, neurally inspired learning mechanisms and the like), these frameworks tend to integrate soft computing features from the subsymbolic world into a symbolic, logic-based framework in order to extend the range of applicability. Additionally, such approaches are often cognitively inspired and often based (at least to a certain extent) on insights from cognitive science.

3 Adaptation from a Symbolic Perspective: An Example

3.1 Heuristic-Driven Theory Projection

In this section, we give informally an example of a symbolic system that allows dynamic adaptations of representations and the learning of new concepts (and theories for these concepts) that are potentially formulated in different languages. Heuristic-Driven Theory Projection (HDTP) is an expressive formalism for the computation of candidates of analogical relations between two given first-order theories (source domain and target domain) [Gust *et al.*, 2006]. Additionally to the computation of an analogical relation the system computes a generalization of the input theories, substitutions in order to recover subtheories of the input theories from the generalization, and the re-representation of input theories (if necessary) [Schwering *et al.*, 2009]. Recently, the framework has been extended to cover also other cognitive mechanisms like concept blending, a mechanism that is important for creativity aspects of cognition [Martinez *et al.*, 2012].

HDTP's computation of an analogical relation identifies the common (structural) aspects of source and target, and exports some of these aspects from the source domain to the target domain. The generalized domain identifies common aspects of the input domains and makes these explicit, i.e. the generalized domain captures the common parts of the input domains. Figure 1 depicts this overall idea using S and T as source and target domains, respectively, and the generalization, G, represents the common parts of S and T. The *analogical transfer* not only associates S and T with each other, but also projects knowledge from S to T, resulting in an enrichment of structure in T.

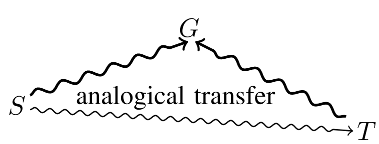

Figure 1: HDTP's overall approach to creating analogies.

The inverse operation of the generalization, namely to reconstruct the input domains from the generalization can be modeled by substitutions and therefore result in specializations of the generalization. Obviously, there are many possibilities to associate entities of two domains resulting in different candidates of analogies. Furthermore, analogies are not correct or incorrect, but more or less psychologically plausible. HDTP takes these features of analogy-making into account by the computation of candidates of analogical relations together with a ranking of these candidates.

Analogical transfer often results in structure enrichment of the target side (via the analogical transfer), which usually corresponds to the addition of new axioms to the target theory, but may also involve the addition of new or the deletion of old first-order symbols. Taking into account the whole process of analogy making, even worse, operations like the replacement of a first-order symbol of arity n by a new first-order symbol of arit m, or a combination of the mentioned operations can occur. Dynamic changes of the underlying signature of a theory are usually not considered in classical logic frameworks, because it is often hard to find a semantics for such dynamical mappings (provided we are not restricting our considerations to well-known conservative expansions or reducts in the model theoretic sense).[3]

Besides the computation of an analogical relation there are application cases in which two conceptual spaces (in our case the input theories for the source and target domains) need not only to be (partially) mapped onto each other, but partially merged in order to create a new conceptual space. In such cases, HDTP uses the computed generalization, the given source and target theories, and the analogical relation between source and target to compute a new conceptual space which is called a blend space [Goguen, 2006].

3.2 Institutions

Analogy making and concept blending as considered in HDTP can be seen as a theory integrating soft computing features into a symbolic framework. Here are some examples of these soft computing features:

- Identification of cross-domain properties and relations that cannot be associated in classical frameworks.
- Adaptation of the underlying input theories (re-representation based on logical deductions) if this is necessary for the computation of better analogies.
- Dynamic transfer of knowledge from the source to the target domain.
- Ranking of candidates by a cost function or an appropriate probability measure.
- Mapping dynamically signatures of underlying domain theories onto each other.

In particular the last issue, namely the dynamic mapping of signatures of logical theories for the analogy-making process is a difficult operation in logical systems. HDTP uses the theory of institutions [Diaconescu, 2008] to specify a semantics for this operation. Figure 2 shows the basic idea of an "institution" in a diagrammatic representation. The theory of institutions allows to represent the elementary dynamics in logical systems if a theory formalized in a language L is mapped to a theory formalized in another language L'. Furthermore, institution theory is an abstract formalism that allows to represent both syntax and semantics in the described dynamic

[3]Compare [Chang and Keisler, 1973] for more information concerning conservative expansions of languages and models of logical theories.

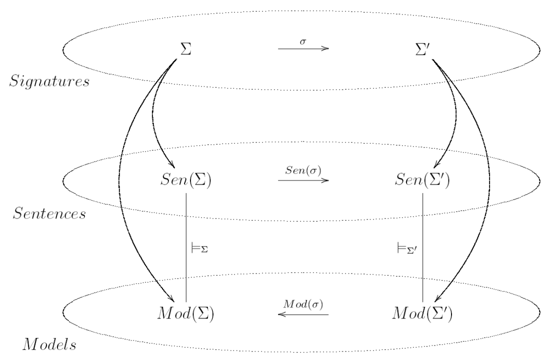

Figure 2: A graphical representation covering the basic idea of an institution.

change from one language to another language. On the signature level a morphism $\sigma : \Sigma \to \Sigma'$ maps a signature Σ to another signature Σ'. Via a functor mapping objects and morphisms in *Signature* to objects and morphisms in *Sentences*, it is possible to induce a morphism $Sen(\sigma)$ from sentences $Sen(\Sigma)$ formulated in Σ to sentences $Sen(\Sigma')$ formulated in Σ'. Similarly, it is possible to induce (via another functor) a morphism $Mod(\sigma)$ from classes of models for $Sen(\Sigma')$ to classes of models of $Sen(\Sigma)$. The important property is the contravariant relation between the morphism between model classes and the morphism between sets of sentences. More precisely, for every $\sigma : \Sigma \to \Sigma'$ every model $\mathfrak{m}' \in Mod(\Sigma'))$ and every $\rho \in Sen(\Sigma)$ it holds:

$$\mathfrak{m}' \models_{\Sigma'} Sen(\rho)(\sigma) \Longleftrightarrow Mod(\sigma)(\mathfrak{m}') \models_{\Sigma} \rho$$

A simple example of an institution can be given in well-known terms of first-order logic (FOL). In this case, the Σ-sentences $\mathbf{Sen}(\Sigma)$ corresponds to the set of all FOL formulas that can be built using symbols from a signature Σ. For each signature Σ the collection $\mathbf{Mod}(\Sigma)$ of all Σ-models corresponds in FOL to the collection of all possible interpretations of symbols from Σ. The Σ-models and Σ-sentences are related by the relation of Σ-satisfaction, which corresponds to the classical model theoretic satisfaction relation in FOL. Institutions theory provides therefore a framework to consider the syntax and the semantics of FOL in one framework.[4]

3.3 Modeling Analogies in Institutions

The described dynamics that can be represented in institution theory is triggered by signature morphisms. In very simple situations, such mappings may suffice to model the analogical relation between a given source and a target domain in an analogical relation. Nevertheless, the general case requires a more general concept of signature morphism, because substitutions specializing terms of the generalized theory may be complex. Such cases are not covered by simple morphism in the category of signatures. Fortunately, institution theory provides such a generalized concept, namely the concept of a generalized Σ.

Definition 1. *For any signature Σ of an institution, and any signature morphisms $\chi_1 : \Sigma \to \Sigma_1$ and $\chi_2 : \Sigma \to \Sigma_2$, a* general Σ-substitution $\psi_{\chi_1:\chi_2}$ *consists of a pair*

$$\langle Sen(\psi), Mod(\psi)\rangle\,,$$

where

- $Sen(\psi) : Sen(\Sigma_1) \to Sen(\Sigma_2)$ *is a function*
- $Mod(\psi) : Mod(\Sigma_2) \to Mod(\Sigma_1)$ *is a functor*

such that both of them preserve Σ, i.e. the following diagrams commute:

$$\begin{tikzcd} Sen(\Sigma_1) \arrow[rr, "Sen(\psi)"] & & Sen(\Sigma_2) \\ & Sen(\Sigma) \arrow[ul, "Sen(\chi_1)"] \arrow[ur, "Sen(\chi_2)"'] & \end{tikzcd}$$

$$\begin{tikzcd} Mod(\Sigma_1) \arrow[dr, "Mod(\chi_1)"'] & & Mod(\Sigma_2) \arrow[ll, "Mod(\psi)"'] \arrow[dl, "Mod(\chi_2)"] \\ & Mod(\Sigma) & \end{tikzcd}$$

and such that the following satisfaction condition holds:

$$Mod(\psi)(\mathfrak{m}_2) \models \rho_1 \quad \textit{if and only if} \quad \mathfrak{m}_2 \models Sen(\psi)(\rho_1)$$

for each Σ_2-model $\mathfrak{m}_2$ and each Σ_1-sentence ρ_1.

General Σ-substitutions extend the idea of a signature morphism. Although, in general there needs to be no mapping on the level of signatures between Σ_1 und Σ_2, most general Σ-substitution considered in practice are induced by some form of signature mapping. Every signature morphism can be seen as general Σ-substitution, and many other mappings, like classical first-order substitutions, second-order substitutions (for FOL), and derived signature morphisms give rise to a general Σ-substitution.

Generalized $\Sigma-$substitutions allow to define formally the concept of an analogy in the sense of HDTP.

Definition 2. *Given two signatures Σ_S and Σ_T over a common signature Σ, an* analogy $\aleph$ *is defined to be a triple $\langle \Sigma_G, \sigma, \tau\rangle$, consisting of a signature Σ_G, and general Σ substitutions σ and τ as indicated in the following diagram:*

$$\begin{tikzcd} & \Sigma_G \arrow[dl, bend right, "\sigma"'] \arrow[dr, bend left, "\tau"] & \\ \Sigma_S & \Sigma \arrow[l] \arrow[r] \arrow[u] & \Sigma_T \end{tikzcd}$$

As a Σ-substitution is defined as a pair of mappings on sentence and model level, every analogy gives rise to the following diagrams:

$$\begin{tikzcd} & Sen(\Sigma_G) \arrow[dl, "Sen(\sigma)"'] \arrow[dr, "Sen(\tau)"] & \\ Sen(\Sigma_S) & Sen(\Sigma) \arrow[l] \arrow[r] \arrow[u] & Sen(\Sigma_T) \end{tikzcd}$$

[4] An alternative way to represent the contrainvariant interplay between syntax and semantics of logical frameworks is channel theory [Barwise and Seligman, 1997].

and

$$
\begin{tikzcd}
& Mod(\Sigma_G) \arrow[d] & \\
Mod(\Sigma_S) \arrow[ur, "Mod(\sigma)"] \arrow[r] & Mod(\Sigma) & Mod(\Sigma_T) \arrow[ul, "Mod(\tau)"'] \arrow[l]
\end{tikzcd}
$$

Furthermore, for every Σ_G-sentence ρ, every Σ_S-model $\mathfrak{m}_S$ and every Σ_T-model $\mathfrak{m}_T$, the following satisfiability conditions hold:

$$
\begin{tikzcd}
\mathbf{Sen}(\sigma)(\rho) \arrow[d, "\Sigma_S"', "\models"] & \rho \arrow[l] \arrow[d, "\Sigma_G"', "\models"] \\
\mathfrak{m}_S \arrow[r, tail] & \mathbf{Mod}(\sigma)(\mathfrak{m}_S)
\end{tikzcd}
\quad \text{iff} \quad \text{and} \quad
\begin{tikzcd}
\rho \arrow[r] \arrow[d, "\Sigma_G"', "\models"] & \mathbf{Sen}(\tau)(\rho) \arrow[d, "\Sigma_T"', "\models"] \\
\mathbf{Mod}(\tau)(\mathfrak{m}_T) & \mathfrak{m}_T \arrow[l, tail]
\end{tikzcd}
$$

In this setting, we can introduce an analogical relation on the level of sentences as well as on the level of models in the intended sense of the computational model. To be more precise, two formulas $s \in \mathbf{Sen}(\Sigma_S)$ and $t \in \mathbf{Sen}(\Sigma_T)$ are in an analogical relation if and only if there is a formula $g \in \mathbf{Sen}(\Sigma_G)$ of the generalized theory such that s and t can be computed from g by applying the respective generalized $\Sigma-$Substitutions. A similar relations holds with respect to the model classes, hence the theory of institutions allows to model rather nicely on the syntactic and semantic level the mappings of different signatures to each other.

4 Conclusions

Neural-symbolic reasoning and learning can be approached from different directions. One direction is to directly attempt a modeling of higher cognitive abilities (like reasoning) with neural or neurally inspired means. Another possibility may be to equip symbolic frameworks with soft computing features. The claim has been made in this paper that currently several researchers follow this second line of research. We exemplified this claim by specifying some soft computing properties of HDTP, a symbolic framework for analogy-making and concept blending.

The convergence of the two worlds as described in Section 2 does not seem to follow a blueprint. It seems to result from needs of the practical application of a certain computational framework to a specific problem domain. For example, analogy-making considered as a cognitive ability, requires in its modeling a non-crisp approach, the possibility to associate aspects of theories locally, and to relate domain theories formulated in potentially different languages to each other. Such features were sketched in Section 3 resulting in an equipment of a symbolic framework with soft computing features. Here is a second example: Ontology repair systems, as mentioned in Section 2, are triggered from a cognitive perspective, simply because the system is intended to model what scientists are doing if they modify or adapt an existing theory to the needs of underlying constraints. Hence, the usage of soft computing features in such frameworks are natural extensions motivated and inspired by the needs of the application. Insofar, it is not surprising that such systems integrate only certain (seemingly isolated) aspects of the symbolic and subsymbolic world into each other, rather than trying to solve the problem in its full generality, i.e. to attempt to integrate all aspects of the two worlds into each other.

It is rather likely that the large number of approaches for integrating locally soft computing features into existing symbolic frameworks will eventually result in a better understanding of frameworks covering aspects of both worlds. We think that the various approaches for a convergence of the symbolic and subsymbolic world should enable researchers to find a broad reservoir of new and inspiring examples such that the next steps for progression with respect to the important question of neural-symbolic reasoning and learning can be anticipated.

References

[Barwise and Seligman, 1997] Jon Barwise and Jerry Seligman. *Information Flow: The Logic of Distributed Systems*, volume 44 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1997.

[Bundy and Chan, 2008] Alan Bundy and Michael Chan. Towards ontology evolution in physics. In *WoLLIC*, pages 98–110, 2008.

[Chan *et al.*, 2010] Michael Chan, Jos Lehmann, and Alan Bundy. Higher-order representation and reasoning for automated ontology evolution. In *KEOD*, pages 84–93, 2010.

[Chang and Keisler, 1973] C. C. Chang and H. J. Keisler. *Model Theory*, volume 73 of *Studies in Logic*. North-Holland Publishing Company, 1973.

[De Raedt, 2008] Luc De Raedt. *Logical and Relational Learning*. Cognitive Technologies. Springer, 2008.

[Diaconescu, 2008] Răzvan Diaconescu. *Institution-independent Model Theory*. Studies in Universal Logic. Birkhäuser, Basel, 2008.

[Garcez *et al.*, 2002] Artur D'Avila Garcez, Krysia Broda, and Dov Gabbay. *Neural-Symbolic Learning Systems: Foundations and Applications*. Springer, 2002.

[Goguen, 2006] Joseph Goguen. Mathematical models of cognitive space and time. In D. Andler, Y. Ogawa, M. Okada, and S. Watanabe, editors, *Reasoning and Cognition: Proc. of the Interdisciplinary Conference on Reasoning and Cognition*, pages 125–128. Keio University Press, 2006.

[Gust *et al.*, 2006] Helmar Gust, Kai-Uwe Kühnberger, and Ute Schmid. Metaphors and heuristic-driven theory projection (hdtp). *Theoretical Computer Science*, 354:98–117, 2006.

[Gust *et al.*, 2007] Helmar Gust, Kai-Uwe Kühnberger, and Peter Geibel. *Perspectives on Neural-Symbolic Integration*, volume 77 of *Studies in Computational Intelligence*, chapter Learning Models of Predicate Logical Theories with Neural Networks Based on Topos Theory, pages 233–264. Springer, 2007.

[Hammer and Hitzler, 2007] Barbara Hammer and Pascal Hitzler, editors. *Perspectives of Neural-Symbolic Integration*, volume 77 of *Studies in Computational Intelligence*. Springer, 2007.

[Hitzler *et al.*, 2004] Pascal Hitzler, Steffen Hölldobler, and Anthony K. Seda. Logic programs and connectionist networks. *Journal of Applied Logic*, 2(3):245–272, 2004.

[Hofstadter and the Fluid Analogies Research Group, 1995] Douglas Hofstadter and the Fluid Analogies Research Group. *Fluid Concepts and Creative Analogies. Computer Models of the Fundamental Mechanisms of Thought*. Basic Books, New York, 1995.

[Martinez *et al.*, 2012] Maricarmen Martinez, Tarek R. Besold, Ahmed Abdel-Fattah, Helmar Gust, Martin Schmidt, Ulf Krumnack, and Kai-Uwe Kühnberger. Theory blending as a framework for creativity in systems for general intelligence. In Pei Wang and Ben Goertzel, editors, *Theoretical Foundations of Artificial General Intelligence*, volume 4 of *Thinking Machines*, pages 219–239. Atlantis Press, Springer, Amsterdam, 2012.

[Ovchinnikova, 2012] Ekaterina Ovchinnikova. *Integration of World Knowledge for Natural Language Understanding*, volume 3 of *Thinking Machines*. Atlantis Press, Springer, Amsterdam, 2012.

[Schwering *et al.*, 2009] Angela Schwering, Ulf Krumnack, Kai-Uwe Kühnberger, and Helmar Gust. Syntactic principles of heuristic-driven theory projection. *Cognitive Systems Research*, 10(3):251–269, 2009.

[Wermter and Sun, 2000] Stefan Wermter and Ron Sun, editors. *Hybrid Neural Systems*. Springer, 2000.