Hybridization Beyond Modularity: A Modern Taxonomy for Fusion and Duality of Data-driven and Knowledge-based Artificial Intelligence

Thomas Schmid[1,2]
[1]Martin Luther University Halle-Wittenberg, Halle (Saale), Germany
[2]Lancaster University Leipzig, Leipzig, Germany

Abstract

Combining data-driven and knowledge-based techniques is a key research direction for developing trustworthy artificial intelligence (AI). Modular hybridization concepts in particular have recently received increased attention in scholary taxonomization concepts, such as the so-called Boxology. At the same time, non-modular hybrid approaches based on fusion and duality have continued to evolve and diversify during the last decade. Building on the canonical typology by McGarry et al, we introduce an updated taxonomy for hybridization approaches that combine data-driven and knowledge-based AI in a non-modular fashion. Notably, we distinguish unified and transformational hybridization into further subcategories and indicate current research trends in these areas. While the original work by McGarry and colleagues was targeting solely neural hybrid architectures, our taxonomy is not limited to neurosymbolic techniques but also includes further hybrid approaches, such as fuzzy-based techniques. Overall, this work seeks to stimulate critical discussion about recent developments and future directions in non-modular hybridization.

1 Introduction

Aiming to combine the complementary strengths of data-driven and knowledge-based artificial intelligence (AI), researchers have been investigating innovative strategies and concepts to hybridize into hybrid systems since the early 1990s. While the intensity of research in this area has been varying since, the underlying motivation of developing reliable, trustworthy, and effective AI systems has continued to be attractive (Martin 2023). With the arrival of large language models (LLMs), approaches for combining machine learning and knowledge engineering have received particular interest among researchers as a strategy to overcome hallucination and other limitations and challenges of LLMs (Boer et al. 2025). Building upon a first wave of knowledge-based AI approaches and a second wave of data-driven AI approaches, the integration of symbolic and subsymbolic methods is today increasingly perceived as an emerging third wave of AI with increasing research activity (van Harmelen 2022; Garcez and Lamb 2023).

A key design principle for such systems is the use and re-use of AI components or modules representing specific functionalities (Schmid 2023b). Analogously to the emergence of design patterns in software engineering (Gamma et al. 1995), the manifold possibilities for composition and interaction have let to comprehensive description systems, such as the so-called Boxology (van Bekkum et al. 2021). In the same spirit, also systems that are not primarily considered as hybrid AI have adapted the concept of integrating complementary AI techniques. Examples for this trend comprise AlphaGo with a combination of Supervised Learning, Reinforcement Learning and Tree Search (Silver et al. 2016) as well as ChatGPT with a combination of self-supervised learning and reinforcement learning from human feedback (Ouyang et al. 2022). As a result, the concept of modularity must today be considered a universal design principle for complex real-world AI systems rather than a specific form of hybridization, as was the case at the beginning of the third wave of AI with 'functional hybrids' (Hilario 1997).

Apart from modularization, a large variety of alternative hybridization strategies have been developed since the early 1990s. This diversity was recognized early on through descriptive classification approaches (Medsker 1994; Hilario 1994) and workshops that highlighted and celebrated the variety of neuro-symbolic systems (Sun and Alexandre 1997). A basic but widely accepted taxonomy published by McGarry et al in 1999 identifies two main categories of hybrid neural systems next to modular hybrid systems (McGarry, Wermter, and MacIntyre 1999): unified hybrid systems, which comprises highly integrated, or fused, complementary techniques, and transformational hybrid systems, which may be employed in more than one way due to a certain level of inheritent duality. Here, we will built upon this high-level classification system, extending it with subcategories and including more recent techniques.

2 Related Work

Most hybridization approaches over the last 30 years have been viewed from a neuro-symbolic perspective or at least been referred to as neuro-symbolic. Notably, the number of papers published using this label has significantly increased over the last decade (Colelough and Regli 2025). This development is reflected by a growing number of reviews and surveys aiming to map the field (Yu et al. 2021; Besold et al. 2021; Dash et al. 2021; Hitzler et al. 2022; Gibaut et al. 2023; Colelough and Regli 2024; Wang, Yang, and

Wu 2024; Liang, Wang, and Tong 2025) as well as by specialized reviews and surveys addressing more specific aspects, such as natural language processing (Panchendrarajan and Zubiaga 2024), knowledge graph reasoning (Cheng et al. 2024), semantic web (Breit et al. 2023) and knowledge-enhanced pretrained language models (Wei et al. 2021). Regarding a taxonomic mapping of the field, perspectives and conclusions vary significantly in the literature, both in breadth and depth. For example, while Medsker found a general distinction between tightly coupled, loosley coupled and fully integrated hybrid systems sufficent to categorize the field (Medsker 1994) decades ago, Kautz more recently proposed six top-level types of neuro-symbolic systems in a widely recognized public talk (Kautz 2020).

Proposed taxonomies and categorizations differ also depending on the context and intention. Focusing in neuro-symbolic integration, Bader and Hitzler suggest to use a three-dimensional mapping employing the categories interrelation, language and usage (Bader and Hitzler 2005). Focusing on explainable AI, integration and composition have been proposed as top-level categories (Calegari, Ciatto, and Omicini 2020), with further distinction into two subcategories for integration (Logic & Numeric vs. Logic & Numeric & Statistic) and two subcategories for composition (Extraction vs. Injection). Focussing on combining LLMs with knowledge representations, Colon-Hernandéz et al. proposed to categorize such systems into three types of so-called knowledge injections (Colon-Hernandez et al. 2021a); with the same intention, de Boer and colleagues proposed to describe such systems using design patterns based on the so-called boxology (Boer et al. 2025)

3 Classifying Non-modular Hybrid Artificial Intelligence

While a substantial increase in research intensity on combining data-driven and knowledge-based AI has lately led to several taxonomies tailored towards specific contexts, we see a need for a novel widely applicable general-purpose taxonomy in this area. To this end, we propose a modernized taxonomy for non-modular hybrid AI built upon the canonical classification scheme proposed by McGarry and colleagues (McGarry, Wermter, and MacIntyre 1999).

More specifically, we adapt the original three-category concept by applying three major modifications:

1. **Focus.** We exclude modular architectures consisting of two or more interacting but structurally independent AI components, as we consider this a universal design concept for any complex AI system and not specific to bipartite combinations of one data-driven and one knowledge-based AI technique (Schmid 2023b).
2. **Scope.** We extend the principles beyond hybrid neural systems and neurosymbolic AI. Notably, hybrid systems employing machine learning techniques other than neural networks can be classified using this taxonomy.
3. **Differentiation.** We extend the foundational categories of unified and transformational systems by introducing an additional layer of subcategories.

Figure 1 provides a hierarchical overview of our modernized taxonomy for non-modular hybrid AI.

Unified Approaches

Originally, McGarry and colleagues introduced the term *Unified Hybrid Systems* for hybridization approaches in which all processing, even functions typically implemented through symbolic rule-based techniques, is carried out within a neural network, making a separate symbolic module unnecessary (McGarry, Wermter, and MacIntyre 1999). Here, we extend this concept to foster data-driven AI, including but not limited to neural networks, and knowledge-based AI, including but limited to rule-based AI.

To this end, we use the term *Unified Hybrid System* for approaches which allow to carry out either typical functionalities of knowledge-based AI by means of data-driven AI or typical functionality of data-driven AI by means of knowledge-based AI. In practice, the case of *Unified Hybrid Sytems* based on data-driven AI aiming to implement typical functionalities of knowledge-based AI will be the more typical case as this scenario enables continuous improvement of the system by learning from experience, which is currently not possible for purely knowledge-based AI.

Other than McGarry and colleagues, we further distinguish the category *Unified Hybrid Systems* into more fine-grained subcategories. Based on key conceptual aspects as well as taking into account the scale of research published around the respectve concepts, we see three major subcategories fitting into the scope of *Unified Hybrid Systems*:

- **Integrative.** *Unified Hybrid Systems* with *integrative* character represent the core concept of this category by integrating the functionality of one distinct AI technique inseparably into another algorithm or at least within component that are not trivially separable. Classic examples for this subcategory include CONSYDERR (Sun 1994), CHCL (Hölldobler and Kurfeß 1992), and SC-NET (Romaniuk and Hall 1993). More recent examples for this category include Neural Logic Machines (Dong et al. 2019), Logical Neural Networks (Riegel et al. 2020), aand Constructivist Machine Learning (Schmid 2023a). Moreover, integration frameworks without specific names have been introduced (Komendantskaya 2007).
- **Fuzzy-based.** *Unified Hybrid Systems* with *fuzzy-based* character combine fuzzy logic with either data-driven or knowledge-based AI. Allowing reasoning with degrees of truth rather than binary true/false values (Zadeh 1965), fuzzy logic provides a powerful interface between these two paradigms. Classic examples for this subcategory include ANFIS (Jang 1993), NEFCLASS (Nauck and Kruse 1996), NEFCON (Nauck and Kruse 1997), and DENFIS (Kasabov and Song 2002). More recent examples for this category include Fuzzy Description Logics (Borgwardt and Peñaloza 2017), DCNFIS (Yeganejou et al. 2023), NASP-T (Machot and Machot 2025), and KANFIS (Yong et al. 2026).
- **Embedding-based.** *Unified Hybrid Systems* with *embedding-based* character represent a considerably

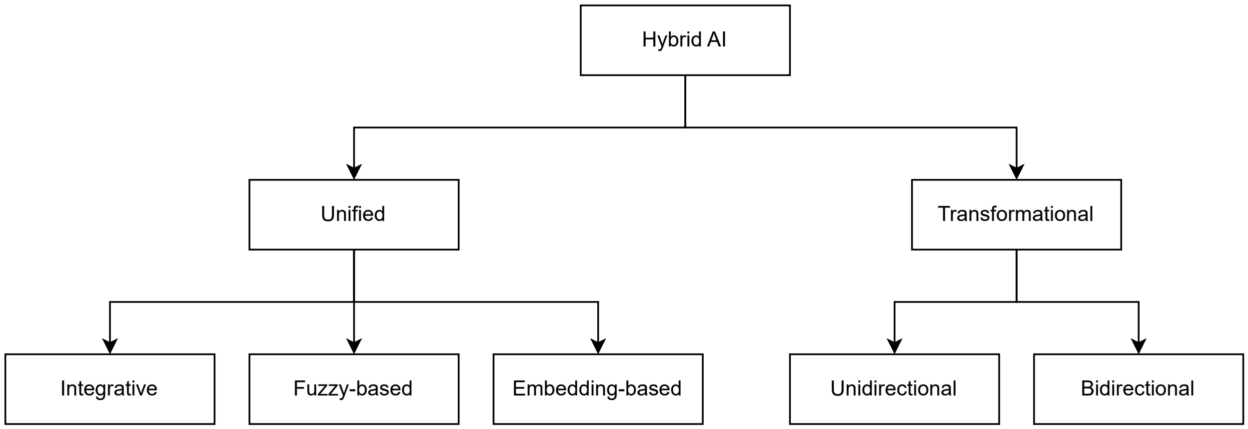

Figure 1: Overview of the proposed taxonomy for fusion and duality of data-driven and knowledge-based artificial intelligence.

dominant trend currently observable in the domain of knowledge-enhanced LLMs. Classic examples for this subcategory are approaches originating from the domain of knowledge base completion (Wang et al. 2015; Wei et al. 2015; Rocktäschel, Singh, and Riedel 2015; Guo et al. 2016). More general examples for this subcategory include Align, Mask, and Select (AMS) (Ye et al. 2019), COMET (Bosselut et al. 2019), KnowBERT (Peters et al. 2019), SemBERT (Zhang et al. 2020), and OSCAR (Goodwin and Demner-Fushman 2019). In the literature, such embedding-based approaches are also often referred to as knowledge injection, which may be further distinguished regarding whether the embedding is injected to the LLM at input level, architecture level, or output level (Colon-Hernandez et al. 2021b).

Transformational Approaches

Originally, McGarry and colleagues introduced the term *Transformational Hybrid Systems* for hybridization approaches that provide the ability of transferring a symbolic representation into a neural representation and vice versa (McGarry, Wermter, and MacIntyre 1999). Here, we extend this concept to foster data-driven AI, including but not limited to neural networks, and knowledge-based AI, including but limited to symbolic AI.

To this end, we use the term *Transformational Hybrid System* for approaches that allow for a transformation between a representation form of data-driven AI, i.e. an embedding, and a representation form of knowledge-based AI, i.e. a knowledge base. In attempts to make deep learning systems more explainable, the transformation of the underlying representations has gained particular attention (Zarlenga, Shams, and Jamnik 2021; Sokol and Flach 2024).

Other than McGarry and colleagues, we further distinguish the category *Transformational Hybrid Systems* into more fine-grained subcategories. Following commonsense logic, such approaches may be distinguished into techniques that allow for a one-way transformation only (for example, from embeddings to symbols) and techniques that allow to transformations in both directions. To this end, we see two major subcategories for *Transformational Hybrid Systems*:

- **Unidirectional.** *Transformational Hybrid Systems* with *unidirectional* character represent the default approach in this category. Classic examples for this subcategory include Tree-based Neural Nets (Ivanova and Kubat 1995) and TREPAN (Craven and Shavlik 1995). More recent examples include RxNCM (Biswas et al. 2017), RENNE (Chakraborty 2024), and DCDL (Burkhardt et al. 2021).
- **Bidirectional.** *Transformational Hybrid Systems* with *bidirectional* character may to some extend be considered a niche in comparison to unidirectional approches. However, several bidirectional approaches can be found in the literature. Classic examples for this subcategory include MACIE (Gallant 1988), Expert Networks (Hruska et al. 1991), Knowledge-Based Artificial Neural Networks (Towell and Shavlik 1994), and Knowledge-Based Conceptual Neural Networks (Fu 2002). More recent examples for this subcategory include Logic Tensor Networks (Serafini and Garcez 2016) and Neural probabilistic logic programming (Manhaeve et al. 2018).

4 Conclusions

With this work, we have introduced a modern taxonomy for fusion and duality of data-driven and knowledge-based AI. This taxonomy reflects key developments of the past decades in hybridization and overcomes inheritent limitations of purely neurosymbolic taxonomies by an extended scope, providing a general-purpose classification system of hybrization strategies beyond modular AI architectures. At the same time, we acknowledge that some hybridization approaches are not explicitly covered by this taxonomy, for example in conversational learning, active dialogue-based learning or other approaches where hybridization is understood as interaction with a human actor (Srivastava, Labutov, and Mitchell 2019). In future work, we aim to extend this proposal into a full taxonomy including and reviewing additional hybridization concepts and hybrid AI systems.

References

Bader, S.; and Hitzler, P. 2005. Dimensions of neural-symbolic integration-a structured survey. *arXiv preprint cs/0511042*.

Besold, T. R.; d'Avila Garcez, A.; Bader, S.; Bowman, H.; Domingos, P.; Hitzler, P.; Kühnberger, K.-U.; Lamb, L. C.; Lima, P. M. V.; de Penning, L.; et al. 2021. Neural-symbolic learning and reasoning: A survey and interpretation 1. In *Neuro-Symbolic Artificial Intelligence: The State of the Art*, 1–51. IOS press.

Biswas, S. K.; Chakraborty, M.; Purkayastha, B.; Roy, P.; and Thounaojam, D. M. 2017. Rule extraction from training data using neural network. *International Journal on Artificial Intelligence Tools*, 26(03): 1750006.

Boer, M. d.; Smit, Q.; Bekkum, M. v.; Meyer-Vitali, A.; and Schmid, T. 2025. Design Patterns for Large Language Model Based Neuro-Symbolic Systems. *Neurosymbolic Artificial Intelligence*, 1: 1–20.

Borgwardt, S.; and Peñaloza, R. 2017. Fuzzy description logics–a survey. In *International Conference on Scalable Uncertainty Management*, 31–45. Springer.

Bosselut, A.; Rashkin, H.; Sap, M.; Malaviya, C.; Celikyilmaz, A.; and Choi, Y. 2019. COMET: Commonsense transformers for automatic knowledge graph construction. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, 4762–4779.

Breit, A.; Waltersdorfer, L.; Ekaputra, F. J.; Sabou, M.; Ekelhart, A.; Iana, A.; Paulheim, H.; Portisch, J.; Revenko, A.; Teije, A. t.; et al. 2023. Combining machine learning and semantic web: A systematic mapping study. *ACM Computing Surveys*.

Burkhardt, S.; Brugger, J.; Wagner, N.; Ahmadi, Z.; Kersting, K.; and Kramer, S. 2021. Rule extraction from binary neural networks with convolutional rules for model validation. *Frontiers in artificial intelligence*, 4: 642263.

Calegari, R.; Ciatto, G.; and Omicini, A. 2020. On the integration of symbolic and sub-symbolic techniques for XAI: A survey. *Intelligenza Artificiale*, 14(1): 7–32.

Chakraborty, M. 2024. Symbolic Interpretation of Trained Neural Network Ensembles. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 32(05): 695–719.

Cheng, K.; Ahmed, N. K.; Rossi, R. A.; Willke, T. L.; and Sun, Y. 2024. Neural–Symbolic Methods for Knowledge Graph Reasoning: A Survey. *ACM Transactions on Knowledge Discovery from Data*, 18(9): 225:1–225:44.

Colelough, B. C.; and Regli, W. 2024. Neuro-Symbolic AI in 2024: A Systematic Review. In *Proceedings of the First International Workshop on Logical Foundations of Neuro-Symbolic AI (LNSAI 2024)*. Jeju, South Korea: Co-located with IJCAI 2024.

Colelough, B. C.; and Regli, W. 2025. Neuro-symbolic AI in 2024: A systematic review. *arXiv preprint arXiv:2501.05435*.

Colon-Hernandez, P.; Havasi, C.; Alonso, J.; Huggins, M.; and Breazeal, C. 2021a. Combining pre-trained language models and structured knowledge. *arXiv preprint arXiv:2101.12294*.

Colon-Hernandez, P.; Havasi, C.; Alonso, J.; Huggins, M.; and Breazeal, C. 2021b. Combining pre-trained language models and structured knowledge. *arXiv preprint arXiv:2101.12294*.

Craven, M.; and Shavlik, J. 1995. Extracting tree-structured representations of trained networks. *Advances in neural information processing systems*, 8.

Dash, T.; Chitlangia, S.; Ahuja, A.; and Srinivasan, A. 2021. Incorporating domain knowledge into deep neural networks. *arXiv preprint arXiv:2103.00180*.

Dong, H.; Mao, J.; Lin, T.; Wang, C.; Li, L.; and Zhou, D. 2019. Neural logic machines. *arXiv preprint arXiv:1904.11694*.

Fu, L.-M. 2002. Knowledge-based connectionism for revising domain theories. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(1): 173–182.

Gallant, S. I. 1988. Connectionist expert systems. *Communications of the ACM*, 31(2): 152–169.

Gamma, E.; Helm, R.; Johnson, R.; and Vlissides, J. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Boston, MA, USA: Addison-Wesley. ISBN 0-201-63361-2.

Garcez, A. d.; and Lamb, L. C. 2023. Neurosymbolic AI: The 3 rd wave. *Artificial Intelligence Review*, 56(11): 12387–12406.

Gibaut, W.; Pereira, L.; Grassiotto, F.; Osorio, A.; Gadioli, E.; Munoz, A.; Gomes, S.; and dos Santos, C. 2023. Neurosymbolic AI and its Taxonomy: a Survey. *arXiv*, abs/2305.08876.

Goodwin, T. R.; and Demner-Fushman, D. 2019. Bridging the knowledge gap: Enhancing question answering with world and domain knowledge. *arXiv preprint arXiv:1910.07429*.

Guo, S.; Wang, Q.; Wang, L.; Wang, B.; and Guo, L. 2016. Jointly embedding knowledge graphs and logical rules. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, 192–202.

Hilario, M. 1994. An Overview of Strategies for Neurosymbolic Integration. In Sun, R.; and Bookman, L., eds., *Computational Architectures Integrating Symbolic and Neural Processes*. Kluwer Academic Publishers.

Hilario, M. 1997. An Overview of Strategies for Neurosymbolic Integration. In Sun, R.; and Alexandre, F., eds., *Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*, 13–35. Hove, UK: Taylor & Francis / Psychology Press.

Hitzler, P.; Eberhart, A.; Ebrahimi, M.; Sarker, M. K.; and Zhou, L. 2022. Neuro-symbolic approaches in artificial intelligence. *National Science Review*, 9(6): nwac035.

Hölldobler, S.; and Kurfeß, F. 1992. CHCL – A connectionist inference system. In Fronhöfer, B.; and Wrightson, G., eds., *Parallelization in Inference Systems*, 318–342. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-47066-3.

Hruska, S.; Dalke, A.; Ferguson, J.; and Lacher, R. 1991. Expert networks in CLIPS. In *NASA. Johnson Space Center, Second CLIPS Conference Proceedings, Volume 2*.

Ivanova, I.; and Kubat, M. 1995. Decision-tree based neural network. In *European Conference on Machine Learning*, 295–298. Springer.

Jang, J.-S. 1993. ANFIS: adaptive-network-based fuzzy inference system. *IEEE transactions on systems, man, and cybernetics*, 23(3): 665–685.

Kasabov, N. K.; and Song, Q. 2002. DENFIS: dynamic evolving neural-fuzzy inference system and its application for time-series prediction. *IEEE transactions on Fuzzy Systems*, 10(2): 144–154.

Kautz, H. A. 2020. The Third AI Summer. AAAI Robert S. Engelmore Memorial Lecture, Thirty-Fourth AAAI Conference on Artificial Intelligence. Presented February 10, 2020.

Komendantskaya, E. 2007. *Learning and deduction in neural networks and logic*. Ph.D. thesis, PhD thesis, Department of Mathematics, University College Cork, Ireland.

Liang, B.; Wang, Y.; and Tong, C. 2025. AI reasoning in deep learning era: From symbolic AI to neural–symbolic AI. *Mathematics*, 13(11): 1707.

Machot, F. A.; and Machot, F. A. 2025. NASP-T: A Fuzzy Neuro-Symbolic Transformer for Logic-Constrained Aviation Safety Report Classification. *arXiv preprint arXiv:2510.05451*.

Manhaeve, R.; Dumancic, S.; Kimmig, A.; Demeester, T.; and De Raedt, L. 2018. Deepproblog: Neural probabilistic logic programming. *Advances in neural information processing systems*, 31.

Martin, A. 2023. AAAI-MAKE 2023: Challenges requiring the combination of machine learning and knowledge engineering. *AI magazine*, 44(2): 204–205.

McGarry, K.; Wermter, S.; and MacIntyre, J. 1999. Hybrid neural systems: from simple coupling to fully integrated neural networks. *Neural Computing Surveys*, 2(1): 62–93.

Medsker, L. R. 1994. *Hybrid Neural Network and Expert Systems*. Boston: Kluwer Academic Publishers.

Nauck, D.; and Kruse, R. 1996. Neuro-fuzzy classification with NEFCLASS. In *Operations Research Proceedings 1995: Selected Papers of the Symposium on Operations Research (SOR'95), Passau, September 13–September 15, 1995*, 294–299. Springer.

Nauck, D.; and Kruse, R. 1997. A neuro-fuzzy method to learn fuzzy classification rules from data. *Fuzzy sets and Systems*, 89(3): 277–288.

Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.

Panchendrarajan, R.; and Zubiaga, A. 2024. Synergizing machine learning & symbolic methods: A survey on hybrid approaches to natural language processing. *Expert Systems with Applications*, 251: 124097.

Peters, M. E.; Neumann, M.; Logan IV, R. L.; Schwartz, R.; Joshi, V.; Singh, S.; and Smith, N. A. 2019. Knowledge enhanced contextual word representations. *arXiv:1909.04164*.

Riegel, R.; Gray, A.; Luus, F.; Khan, N.; Makondo, N.; Akhalwaya, I. Y.; Qian, H.; Fagin, R.; Barahona, F.; Sharma, U.; et al. 2020. Logical neural networks. *arXiv preprint arXiv:2006.13155*.

Rocktäschel, T.; Singh, S.; and Riedel, S. 2015. Injecting logical background knowledge into embeddings for relation extraction. In *Proceedings of the 2015 conference of the north American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 1119–1129.

Romaniuk, S. G.; and Hall, L. O. 1993. SC-net: a hybrid connectionist, symbolic system. *Information sciences*, 71(3): 223–268.

Schmid, T. 2023a. *Constructivist Machine Learning*. Frontiers in Artificial Intelligence and Applications. Netherlands: IOS Press. ISBN 9781643684062.

Schmid, T. 2023b. A Systematic and Efficient Approach to the Design of Modular Hybrid AI Systems. In Martin, A.; Hinkelmann, K.; Fill, H.-G.; Gerber, A.; Lenat, D.; Stolle, R.; and van Harmelen, F., eds., *Proceedings of the AAAI 2023 Spring Symposium on Challenges Requiring the Combination of Machine Learning and Knowledge Engineering (AAAI-MAKE 2023)*.

Serafini, L.; and Garcez, A. d. 2016. Logic tensor networks: Deep learning and logical reasoning from data and knowledge. *arXiv preprint arXiv:1606.04422*.

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.

Sokol, K.; and Flach, P. 2024. Interpretable representations in explainable AI: from theory to practice. *Data Mining and Knowledge Discovery*, 38: 3102–3140.

Srivastava, S.; Labutov, I.; and Mitchell, T. 2019. Learning to Ask for Conversational Machine Learning. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 4164–4174. Hong Kong, China: Association for Computational Linguistics.

Sun, R. 1994. CONSYDERR: a two-level hybrid architecture for structuring knowledge for commonsense reasoning. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, volume 3, 1475–1480. IEEE.

Sun, R.; and Alexandre, F. 1997. *Connectionist-Symbolic Integration: From Unified to Hybrid Approaches*. Psychology Press, 1st edition.

Towell, G. G.; and Shavlik, J. W. 1994. Knowledge-based artificial neural networks. *Artificial intelligence*, 70(1-2): 119–165.

van Bekkum, M.; de Boer, M.; van Harmelen, F.; Meyer-Vitali, A.; and Teije, A. t. 2021. Modular design patterns for hybrid learning and reasoning systems: a taxonomy, patterns and use cases. *Applied Intelligence*, 51(9): 6528–6546.

van Harmelen, F. 2022. Preface: The 3rd AI wave is coming, and it needs a theory. In Hitzler, P.; and Sarker, M. K., eds., *Neuro-Symbolic Artificial Intelligence: The State of the Art*, V–VII. IOS Press.

Wang, Q.; Wang, B.; Guo, L.; et al. 2015. Knowledge Base Completion Using Embeddings and Rules. In *IJCAI*, 1859–1866.

Wang, W.; Yang, Y.; and Wu, F. 2024. Towards Data-and Knowledge-Driven Artificial Intelligence: A Survey on Neuro-Symbolic Computing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(2): 878–899.

Wei, X.; Wang, S.; Zhang, D.; Bhatia, P.; and Arnold, A. 2021. Knowledge enhanced pretrained language models: A compreshensive survey. *arXiv:2110.08455*.

Wei, Z.; Zhao, J.; Liu, K.; Qi, Z.; Sun, Z.; and Tian, G. 2015. Large-scale knowledge base completion: Inferring via grounding network sampling over selected instances. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, 1331–1340.

Ye, Z.-X.; Chen, Q.; Wang, W.; and Ling, Z.-H. 2019. Align, mask and select: A simple method for incorporating commonsense knowledge into language representation models. *arXiv preprint arXiv:1908.06725*.

Yeganejou, M.; Honari, K.; Kluzinski, R.; Dick, S.; Lipsett, M.; and Miller, J. 2023. DCNFIS: Deep convolutional neuro-fuzzy inference system. *arXiv preprint arXiv:2308.06378*.

Yong, B.; Pei, H.; Shen, J.; Li, H.; Zhou, Q.; and Su, Z. 2026. KANFIS A Neuro-Symbolic Framework for Interpretable and Uncertainty-Aware Learning. *arXiv preprint arXiv:2602.03034*.

Yu, D.; Yang, B.; Liu, D.; and Wang, H. 2021. A Survey on Neural-symbolic Learning Systems. .

Zadeh, L. A. 1965. Fuzzy sets. *Information and control*, 8(3): 338–353.

Zarlenga, M. E.; Shams, Z.; and Jamnik, M. 2021. Efficient decompositional rule extraction for deep neural networks. *arXiv preprint arXiv:2111.12628*.

Zhang, Z.; Wu, Y.; Zhao, H.; Li, Z.; Zhang, S.; Zhou, X.; and Zhou, X. 2020. Semantics-aware BERT for language understanding. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 9628–9635.