

Design Patterns for Large Language Model Based Neuro-Symbolic Systems

Neurosymbolic Artificial Intelligence

Volume 1: 1–20

© The Author(s) 2025

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/29498732251377499

journals.sagepub.com/home/nai



Maaïke de Boer^{1*} , Quirine Smit^{1*} , Michael van Bakkum¹ , André Meyer-Vitali² and Thomas Schmid^{3,4}

Abstract

Large language models (LLMs) have been a dominating trend in artificial intelligence (AI) in the past years. At the same time, neuro-symbolic systems employing LLMs have also received increasing interest due to their advantages over purely statistical generative models: They can make explicit use of expert knowledge and can be understood and inspected by humans thus providing explainability. However, with an increasing variety of approaches, it is currently difficult to compare the different ways in which designing, training, fine-tuning, and applying such approaches take place. In this work, we use and extend the modular design patterns for hybrid learning and reasoning systems and the Boxology language of van Bakkum et al. for this purpose. These patterns provide a general language to describe, compare, and understand the different architectures and methods used for LLM-based neuro-symbolic systems. The primary goal of this work is to support a better understanding of specific classes of such systems, namely LLM-based models that are used in conjunction with knowledge-based (symbolic) systems. In order to demonstrate the usefulness of this approach, we explore existing LLM-based neuro-symbolic architectures and approaches, as well as use cases for these design patterns.

Keywords

design patterns, neuro-symbolic AI, generative models, large language models

Received: September 28, 2024; accepted: July 22, 2025

Editors: Marta Sabou, Vienna University of Economics and Business, Austria; Raghava Mutharaju, IIIT Delhi, India; Frank van Harmelen, Vrije Universiteit Amsterdam, Netherlands

Solicited reviews: Jacopo de Berardinis, University of Liverpool, UK; Three anonymous reviewers

1 Introduction

A major driving force for the origin and growth of the field of artificial intelligence (AI) has been the goal of ‘making computers solve really difficult problems’ (Minsky, 1961). Today, many contemporary observers agree that AI has taken a leap in recent years (Savage, 2020) and has now reached a level of capacity and productivity that was unprecedented in previous decades (OECD, 2023). Since 2010, AI systems have reached a close-to-human or even superior-to-human

¹Department of Data Science, TNO, The Hague, The Netherlands

²Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI), Saarbrücken, Germany

³Medical Faculty, Martin Luther University Halle-Wittenberg, Halle (Saale), Germany

⁴Leipzig School of Computing and Communications, Lancaster University, Leipzig, Germany

*These authors have contributed equally for this work.

Corresponding Authors:

Maaïke de Boer, Department of Data Science, TNO, The Hague, The Netherlands.

Email: maaïke.deboer@tno.nl

Quirine Smit, Department of Data Science, TNO, The Hague, The Netherlands.

Email: quirine.smit@tno.nl



Creative Commons CC BY: This article is distributed under the terms of the Creative Commons Attribution 4.0 License

(<https://creativecommons.org/licenses/by/4.0/>) which permits any use, reproduction and distribution of the work without further

permission provided the original work is attributed as specified on the SAGE and Open Access page

(<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

level in many computer vision tasks (Khanday & Sofi, 2021). Since 2020, an increasing number of AI systems have been introduced that are able to successfully complete complex text generation tasks in natural language processing (NLP) (Min et al., 2023), such as text summarisation, translation, and question answering. More recently, the concept of generation has even been extended to multi-modal approaches involving, for example, text input and image output (Betker et al., 2023; Jin et al., 2024; Rombach et al., 2021; Wang et al., 2024a; Zhang et al., 2024). Many of these systems have demonstrated NLP capabilities at a level very close-to-human capabilities (Zhong et al., 2023).

These developments are largely attributed to advances in deep learning techniques, in particular in the form of generative AI and Large language models (LLMs). A wealth of different LLM models have been and are being developed and published, both open-source and proprietary (Chen et al., 2023; Dubey et al., 2024; Kukreja et al., 2024; Minaee et al., 2024). The key technology most current LLMs use is the transformer architecture. The original transformer architecture published by Vaswani et al. (2017) proposed using two interacting models, an encoder and a decoder. These can be trained end-to-end (such as *flan-T5* Chung et al., 2022). Alternatively, architectures have been proposed using encoder-only (BERT Devlin, 2018) or decoder-only (GPT Brown et al., 2020, BLOOMZ Muennighoff et al., 2022, PaLM Chowdhery et al., 2023) models. As only a few LLMs based on other architectures have been proposed Peng et al. (2023) and Beck et al. (2024), in this paper, we focus on transformer-based LLMs and consider encoder-only, decoder-only, and encoder-decoder systems to be possible types of LLMs.

These different categories of transformers provide different advantages and disadvantages, depending on the intended scenario. Encoder-only transformers, such as BERT (Devlin, 2018), specialise in contextual encoding, often named base models. They use context to encode input sentences and represent them as machine-interpretable representations, such as vector representations. Decoder-only systems are complementary to the encoder-only paradigm, but structurally different (Min et al., 2023). A decoder-only system decodes the input data directly, without being transformed into a higher and more abstract representation to the desired representation (text, images, or otherwise). Examples of this are generative models from the GPT family (Brown et al., 2020). Decoder-only architectures can be further divided into causal decoder architectures and prefix decoder architectures. Causal decoder architectures, such as GPT (Brown et al., 2020; Radford et al., 2019) and BLOOMZ (Muennighoff et al., 2022), use only unidirectional attention to the input sequence by using a specific mask. Prefix decoder architectures, such as PaLM (Chowdhery et al., 2023), use bidirectional attention for tokens in the prefix while maintaining unidirectional attention to generate subsequent tokens.

Despite the many impressive achievements and capabilities of many LLMs, a wide variety of challenges remain for purely statistical LLMs (Kaddour et al., 2023). This includes not only substantial costs for training (Schwartz et al., 2020) and inference (Samsi et al., 2023), but also the infamous phenomenon of hallucination: Situations where a trained LLM generates outputs presented as plausible or authoritative information that are factually incorrect, nonsensical, or unfaithful to the input or context (Ji et al., 2023). In general, this phenomenon is attributed to the lack of training data for certain outputs, often indicating a lack of domain-specific knowledge (Zhao et al., 2023). While this may in part be corrected by retrieval-augmented generation (RAG) (Gao et al., 2023) and fine-tuning (Zhang et al., 2023), they introduce novel challenges by themselves; fine-tuning, for example, may well raise significant additional training costs and may lead to catastrophic forgetting. A challenge in using RAG is that the performance is highly dependent on the quality and accuracy of the information retrieved. Moreover, both hallucinations and lack of domain-specific knowledge are often hard to identify or contextualise due to the inherent lack of explainability and interpretability of LLMs (Pan et al., 2024; Zhao et al., 2024). It is therefore fair to state that current LLM-based systems and applications lack a sufficient degree of trustworthiness (Huang et al., 2024; Lin et al., 2024), rendering them unusable where reliable output is essential (e.g., in scientific discovery Schmidt et al., 2024).

In response to these challenges, a variety of novel neuro-symbolic approaches to LLM-based AI systems have recently emerged (Hitzler et al., 2022; Wei et al., 2021). Due to the quantity and diversity of emerging generative techniques, it becomes increasingly challenging to keep track of the ever-growing variety of models with different LLM architectures and capabilities (Colelough & Regli, 2024). This becomes even more challenging with the growing diversity in combining LLMs with symbolic AI techniques using different strategies on different architectural levels and training stages (Amador-Domínguez et al., 2024; De Raedt et al., 2020). A practical solution to tackle the issue of analysing and understanding these approaches in a systematic way is to apply a high-level conceptual framework to discuss, compare, configure, and combine different models. Such a framework is provided by the Boxology, introduced by Van Harmelen and Ten Teije (2019) in 2019. The Boxology was extended in 2021 van Bekkum et al. (2021) by providing a taxonomically organised vocabulary to describe both processes and data structures used in hybrid systems. The Boxology represents a flexible and widely applicable framework for representing AI design patterns.

In this paper, which is an extension of our previous work on LLM-based neuro-symbolic systems (de Boer et al., 2024), we propose to use and extend Boxology to gain insight into a variety of LLMs, specifically on LLMs used in a neuro-symbolic approach. To this end, this paper provides two contributions: Firstly, we propose novel design patterns as

an extension of the current Boxology to promote transparency and trustworthiness in system design, by providing interpretable, high-level component descriptions of LLM-based neuro-symbolic systems. Our modular approach supports new architectures and engineering approaches to LLM-based systems. Secondly, we test validity and usefulness of the Boxology and our extensions in this field on example architectures and applications, such as ChatGPT, KnowGL, GENOME and Logic-LM.

The remainder of the paper is organised as follows. In the next section, we give a more detailed overview of the related work regarding LLMs and LLM-based neuro-symbolic systems. In the third section, we propose to extend the Boxology by three novel basic patterns in order to be able to handle LLMs, and in the fourth section we explain several compositional design patterns in this field. In section 5, we dive into specific applications and tasks in which LLMs, specifically in neuro-symbolic systems, are used. We conclude with a discussion (section 6) and a conclusion summarising our key findings and outlining future work in section 7.

2 Related Work

2.1 Neuro-Symbolic Systems and Design Patterns

For a long time, the field of AI developers had been shaped by an opposition between ‘imitators of the mind’ favouring symbolic AI approaches and ‘imitators of the brain’, favouring statistical approaches, such as artificial neural networks (Cordeschi, 2007). Today, however, combining the complementary strengths and weaknesses of symbolic and statistical approaches is considered crucial for the creation of reliable, trustworthy, and effective AI systems (Martin, 2023). After a first wave of symbolic approaches and a second wave of statistical approaches, the combination of symbolic and statistical approaches is therefore anticipated as the latest, third wave of AI (Garcez & Lamb, 2023; van Harmelen, 2022). Starting in the early 1990s, neuro-symbolic AI has been emerging with a large variety of diverse approaches. This characteristic has been recognised early with descriptive classification approaches (Hilario, 1994; Medsker, 1994) and workshops that showcase and celebrate the diversity of neuro-symbolic systems (Sun & Alexandre, 1997). A basic yet widely recognised and applicable taxonomy published by McGarry et al. in 1999 differentiates three main classes of hybrid neural systems (McGarry et al., 1999): Unified hybrid systems, translational hybrid systems and modular hybrid systems.

Unified and translational hybrid systems as defined by McGarry have been a popular area of study within the neuro-symbolic AI community during the last two decades (Besold et al., 2021; Yu et al., 2021). At the same time, modular hybrid systems have also gained wider interest due to its relevance for designing modern AI systems under industrial conditions (Schmid, 2023). In general, the design paradigm of modularisation is characterised and motivated by the possibility to reuse and/or reorganise modules. Identifying and characterising design patterns for modular approaches of neuro-symbolic has led to the development of the so-called Boxology framework described below (van Bekkum et al., 2021; Van Harmelen & Ten Teije, 2019). This framework has been used and extended in different ways, such as the formalisation of the notions from the Boxology and implementation in the heterogeneous tool set (Mossakowski, 2022), the extension of the Boxology for (teams of) actors (Meyer-Vitali et al., 2021), the characterisation of emerging data-driven knowledge engineering trends (Sabou et al., 2024), and the systematic study of nearly 500 papers published in the past decade in the area of Semantic Web Machine Learning (Breit et al., 2023). While the original Boxology framework focussed on architectural aspects of hybrid learning and reasoning, it has been pointed out the necessity of including representations of human involvement into such patterns (Meyer-Vitali et al., 2021; Witschel et al., 2020). Despite the usefulness and manifold advantages of the above-mentioned traditional taxonomies and studies, however, it must be stated that none of these are able to foresee and fully capture the characteristics of the latest generation of neural systems (de Boer et al., 2024), in particular the broad capabilities and the required extraordinary architectural complexity of LLMs.

A recent approach to describe modular neuro-symbolic AI systems is the ontological visual framework termed EASY-AI, which uses semantically enhanced symbols to represent the components and architectures of the AI system (Ellis et al., 2024a). EASY-AI aims to provide a standardised symbolic language for conveying the structure, purpose, and characteristics of AI systems. The approach presents the logical formalisms underpinning this visual framework, with the objective of enhancing the comprehensibility and understandability of AI system behaviours. Recently, this framework has also been provided with an initial implementation named SNOOP-AI (Ellis et al., 2024b). This framework and implementation could be used in the implementation of the design patterns, as it can provide a formal conceptual foundation for the design patterns that allows formal reasoning over (compositions of) its elements. To the best of our knowledge, specific LLM-based use cases have not been tested using formalisation and implementation yet.

2.2 LLM-based Neuro-Symbolic Systems

While older frameworks for neuro-symbolic AI have not been designed to reflect the latest generation of complex large-scale neural components in the first place, some alternative approaches for combining LLMs and knowledge-based components have been designed from scratch. Colon-Hernández et al. were among the first to do so by identifying three different categories of so-called knowledge injections (Colon-Hernandez et al., 2021): Approaches to modify the architecture of LLM by adding additional layers that integrate knowledge with contextual representations or by modifying existing layers (termed architectural injections) are distinguished from approaches aiming to modify either the structure of the input or the data selected to be fed into the LLM (referred to as input injection) and approaches to change either the output structure or the losses that were used in the base model in some way to incorporate knowledge (termed output injection). While providing new insights and perspectives for LLM-based neuro-symbolic systems, however, this injection-oriented approach is centred largely around modifications of LLMs, leaving out the modular perspective of integrating equally relevant neural and symbolic components.

In light of the popularity and success of knowledge graphs (KGs) in recent years (Ji et al., 2021), it is not surprising that this symbolic technique is a prime candidate for many researchers thriving the enhance LLMs. To this end, Agrawal et al. have suggested to distinguish between LLM-based systems with knowledge-aware inference (KG-augmented retrieval, KG-augmented reasoning, or KG-controlled generation), knowledge-aware training (pre-training or fine-tuning), and knowledge-aware validation (Agrawal et al., 2023). While their taxonomy represents a rather empirical high-level categorisation, an alternative approach by Pan et al. (2024) provided a more differentiated approach where a distinction is made between KG-enhanced LLMs, LLM-augmented KGs and synergised LLMs + KGs. For KG-enhanced LLMs, two primary approaches have been explored: Incorporation during the pre-training stage to facilitate knowledge acquisition and utilisation during the inference stage to improve access to domain-specific information. Additionally, KGs have been employed post hoc to augment the interpretability of LLMs, elucidating both factual content and reasoning processes. In order to augment KGs, LLMs have been employed as text encoders to enrich KG representations and extract relations and entities from the original corpora. Recent studies have focussed on designing KG prompts that effectively convert structural KGs into LLM-comprehensible formats, allowing direct application of LLMs to KG-related tasks such as completion and reasoning. Moreover, the authors have proposed considering the effects and concepts of synergised LLM + KG with respect to four layers: (1) Data, (2) Synergised Model, (3) Technique, and (4) Application. We will loosely use the categorisation of this paper in our exploration of different LLM-based neuro-symbolic systems.

3 Extending the Boxology Framework with Novel Elements

We base our paper on the previous work of van Bekkum and colleagues (van Bekkum et al., 2021), in which a taxonomically organised vocabulary is provided to describe both processes and data structures used in neuro-symbolic systems. The highest level of this taxonomy contains instances, models, processes, and actors, which may be described as follows:

Instances: The two main classes of instances are data and symbols. *Symbols* are defined as having a designation to an object, class, or relation in the world that can be atomic or complex, and when a new symbol is created from another symbol and a system of operations, it should have a designation. Examples of symbols are labels (short descriptions), relations (connections between data items, such as triples), and traces (records of data and events). *Data* is defined as not symbolic. Examples are numbers, texts, tensors, or streams.

Models: Models are descriptions of entities and their relationships, which can be statistical or semantic. *Statistical* models represent dependencies between statistical variables, such as LLMs or Bayesian Networks. *Semantic* models specify concepts, attributes, and relationships to represent the implicit meaning of symbols, such as ontologies, taxonomies, KGs, or rule bases.

Processes: Processes are operations on instances and models. Three types of processes are defined: Generation, transformation, and inference. *Generation* can be performed using, for example, the training of a model or by knowledge engineering. *Transformation* is the transformation of data, for example, from aKG to vector space. *Inference* can be inductive or deductive, in which induction generalises instances and deduction reaches conclusions on specific instances, such as classification.

Actors: Actors can be humans, (software) agents, or robots (physically embedded agents). Meyer-Vitali et al. (2021) extended the original paper with a definition of teams of actors in the Boxology.

In addition to vocabulary, visual language is defined in van Bekkum et al. (2021), as an extension of Van Harmelen and Ten Teije (2019). The visual language consists of rectangular boxes (instances), hexagonal boxes (models), ovals (processes) and triangles (actors), and untyped arrows between them. Within the boxes the concept will be noted by each level in the vocabulary using colon-separation from most generic to most-specific, for example a neural network will be

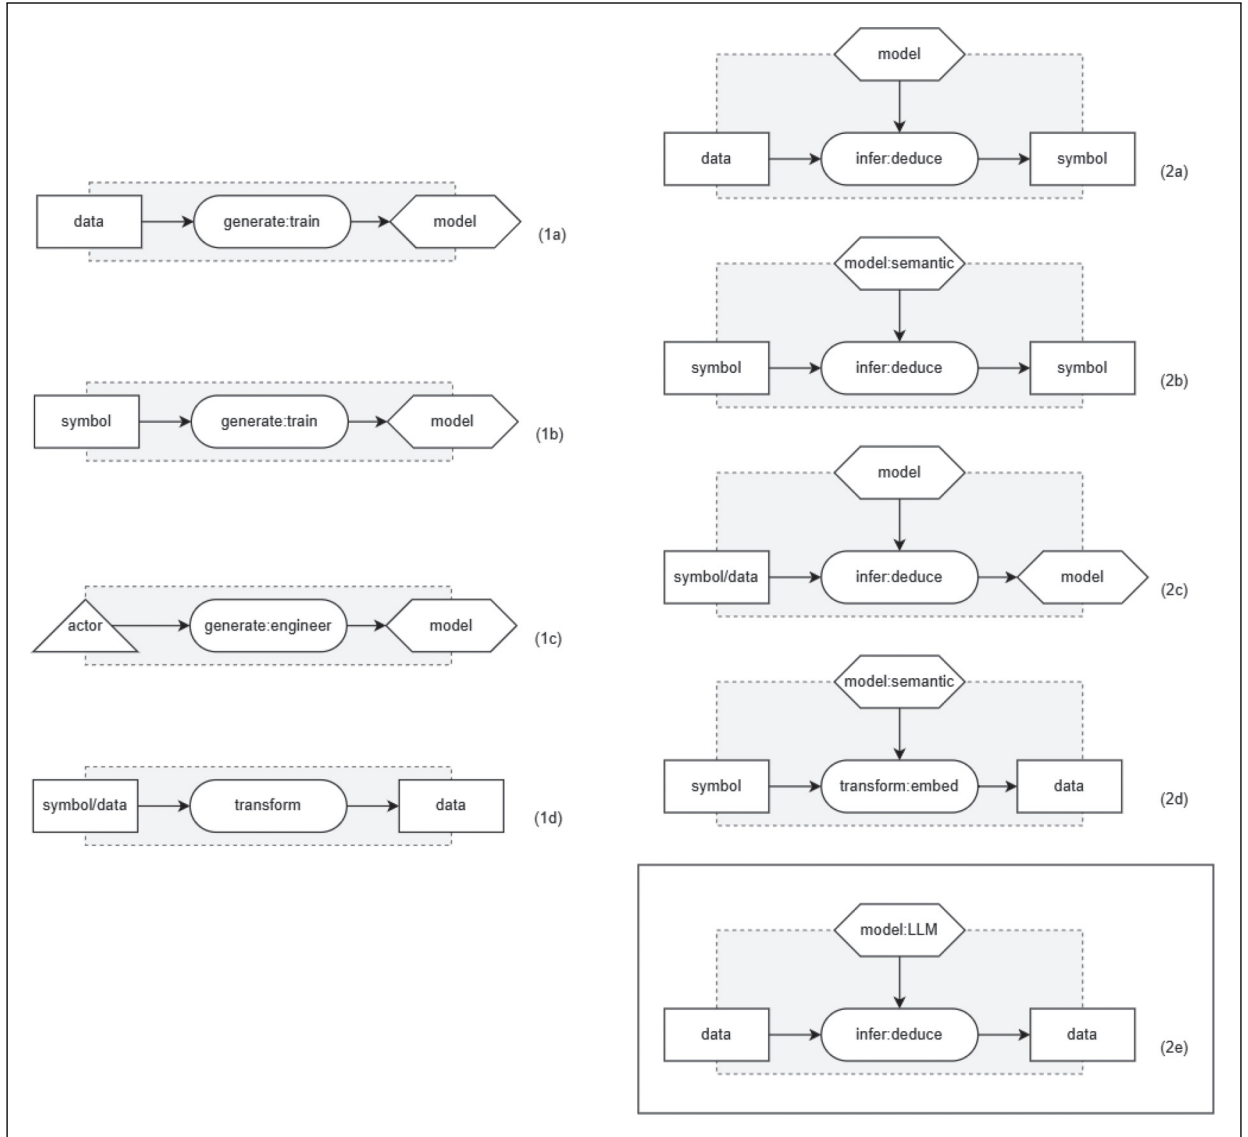


Figure 1. All elementary design patterns, including novel addition 2e. Patterns 1a to 1c allow for model generation, 1d for transforming data and patterns 2a–2e allow for model use.

model:stat:NN. In the figures we use the notation of most generic and more specific only to improve readability, so model:NN.

3.1 Introducing a New Elementary Pattern

The classic Boxology framework van Bekkum et al. (2021) is based on eight elementary patterns (cf. Figure 1). The elementary patterns 1a–1d are elementary to generate a model, in particular for both statistical as well as semantical models. For instance, Pattern 1a shows how to train a statistical model (such as a neural network) using data (such as text or images). Pattern 1b shows how to create a semantic model (such as a KG) using symbols (such as triples). The elementary patterns 2a–2d are patterns describing how to use a model. Pattern 2a, for example, shows a statistical model (such as a neural network) being used to deduce symbols, i.e. for a classification task. Pattern 2b depicts the application of a semantic model, for example when using it for reasoning.

When seeking to apply any of the eight elementary patterns of the classic Boxology framework to represent LLMs, however, it becomes apparent that such generative AI approaches are not adequately represented in these elementary patterns. While the main characteristic of generative AI systems is the ability to output new data for given input data,

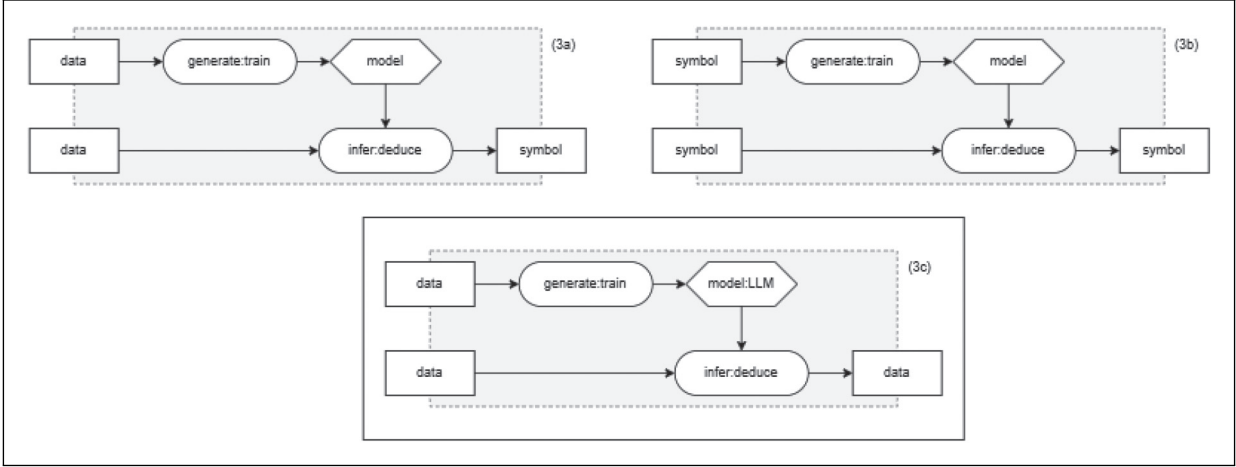


Figure 2. Compositional design patterns, including novel addition 3c made by combining elementary pattern 1a and 2e. Patterns 3a and 3b visualise the patterns for full learning and prediction tasks from data/symbols.

in the existing elementary patterns only symbols or models can be inferred. Therefore, we propose to extend the eight elementary patterns (Figure 1, 1(a)–(d) and 2(a)–(d)) by introducing a new additional elementary pattern 2e (Figure 1). In contrast to the existing elementary patterns, this novel pattern allows to represent LLM-based neuro-symbolical systems since this pattern represents a model that can infer new data from data. This new data can be an image, video, or text, depending on the type of model. For example, with GPT (Brown et al., 2020), LLaMa (Dubey et al., 2024) and similar generative text models, new text is generated based on given input text.

While we focus on textual models in this paper, it is worth mentioning that the pattern proposed in this section abstracts from the specifics of the type of data. The new elementary pattern 2e is thus transferable to other generative models and data types and applies also for example to image generation models (Betker et al., 2023; Jin et al., 2024; Rombach et al., 2021; Wang et al., 2024a; Zhang et al., 2024), which can generate image data from text data. Specifically this would mean that the type of input data would be `data:text` and the type of output data would be `data:image`.

3.2 Introducing a New Compositional Pattern

One of the key characteristics of the Boxology framework is its modularity, allowing the combination and reuse of elementary patterns in compositional patterns. Van Bekkum et al. describe, for example, the two compositional patterns 3a and 3b depicted in Figure 2 (van Bekkum et al., 2021). For the compositional pattern 3a, the two elementary patterns 1a and 2a (Figure 1) are combined. The resulting compositional pattern describes a basic structure for a (statistical) machine learning model depicting the training (creation of the model) and testing or application phase (application of the model on new data). Similarly, the compositional pattern 3b is created from combining two elementary patterns, allowing to depict a basic structure for a semantic model.

When seeking to use any of the two compositional patterns 3a and 3b as a combined representation for training and application of an LLM, however, it becomes apparent that—just as with elementary patterns for LLM-based training and inference—the compositional patterns of the classic Boxology framework are not well-suited for this. Thus, we introduce the new compositional pattern 3c (Figure 2) for the combined training and application of an LLM. This additional pattern is built from a combination of the two elementary patterns 1a and 2e (Figure 1). Similar to the novel elementary pattern 2e, this novel compositional is not limited to text-based data only, but can also be applied to any type of data, including image data, audio data, and multimodal data.

4 LLM-based Neuro-Symbolic Design Patterns

Current LLM-based neuro-symbolic systems often either use an LLM followed by a semantic model, or a semantic model followed by an LLM, or a combination of two models in parallel of which the output is fused. In this section, we propose compositional design patterns for these different types of system. We loosely follow the categorisation of Pan et al. (2024). We divide the section into training and application phases, as the patterns for these phases are distinct.

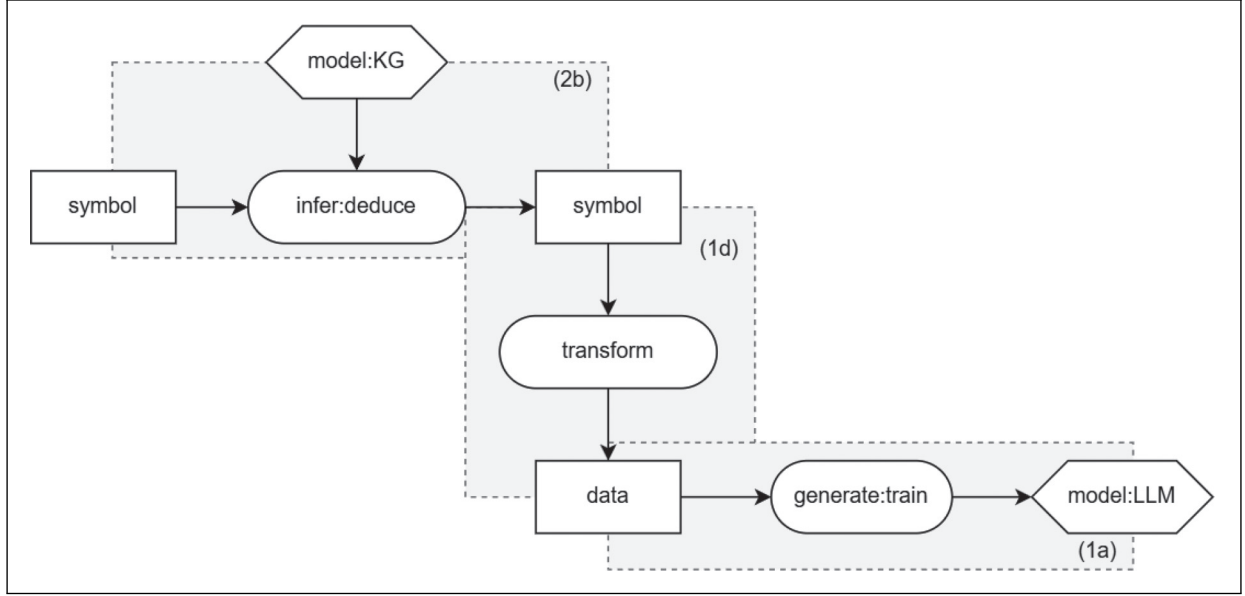


Figure 3. KG-enhanced LLMs in training. KG: knowledge graph; LLM: large language model.

4.1 LLM-based Neuro-Symbolic Design Patterns in Training

Generative neuro-symbolic systems can use semantic models in the training of an LLM or use an LLM to create a semantic model, or can be used in synergy to create a model. In the following subsections, we will describe the different patterns in more detail.

4.1.1 KG-enhanced LLMs. KGs can be used to enhance LLMs in training, for example by influencing the training data. An example of this is shown in the design pattern in Figure 3. Here, a KG is used to infer symbols (pattern 2b). These symbols are then changed into data (pattern 1d). This data is then used to train the LLM (pattern 1a). This depiction can be used to represent (Li et al., 2022; Rosset et al., 2020; Shen et al., 2020; Xiong et al., 2019), for example when a KG is used when masking the data to improve the training of the LLM. For example, in GLM (Shen et al., 2020) the masking probability is higher for concepts which are close together in the KG. On the other hand, SKEP (Tian et al., 2020) uses a KG to identify words of high sentiment and gives them a higher masking probability.

4.1.2 LLM-augmented KGs. LLMs can be used to enhance KGs as KGs might be incomplete and textual information is not integrated in the embedding itself. It can be represented in Boxology as presented in Figure 4. Similar to the KG-enhanced LLMs in training, new data is inferred by an LLM (pattern 2e). These are then transformed to symbols (1d) and used to create or add on to a KG model (1a). For example, Nayyeri et al. (2023) generate representations on different levels such as sentence and document using LLMs and Huang et al. (2022) create multi-modal embeddings. Models following this structure are often used for tasks such as LLM-augmented KG completion and construction, including Named Entity Recognition, Coreference Resolution, and Relation Extraction. For example, KG-BERT, MLT-KGC, and PKGC use LLMs for the completion of a KG (Kim et al., 2020; Lv et al., 2022; Yao et al., 2019). They use the LLM output to predict the relation between new entities and existing ones. Yan et al. (2021) uses LLMs to aid in Named Entity Recognition, Cattani et al. (2021); Joshi et al. (2020) for Coreference Resolution, and Park and Kim (2021); Shi and Lin (2019) for Relation Extraction.

4.1.3 Synergised LLMs and KGs. One of the ways in which LLMs and KGs are synergised in training is using an LLM for joint text and KG embedding or representation. Figure 5 shows the Boxology representation of these types of systems. The symbolic triples are transformed into text (pattern 1d), which is then combined with other text to integrate both the graph structure and the textual information into the embedding simultaneously and trained to create a model (pattern 1a). For example, kNN-KGE sees entities as special tokens and incorporates them into sentences as input for the LLM (Wang et al., 2023b). LMKE has a similar system structure but applies a different learning method to improve the learnt embeddings (Wang et al., 2022). LambdaKG improves the representation of the graph structure by including neighbouring entities in the input sentence (Xie et al., 2022). KEPLER, JointGT and DRAGON use a unified model for the knowledge embedding

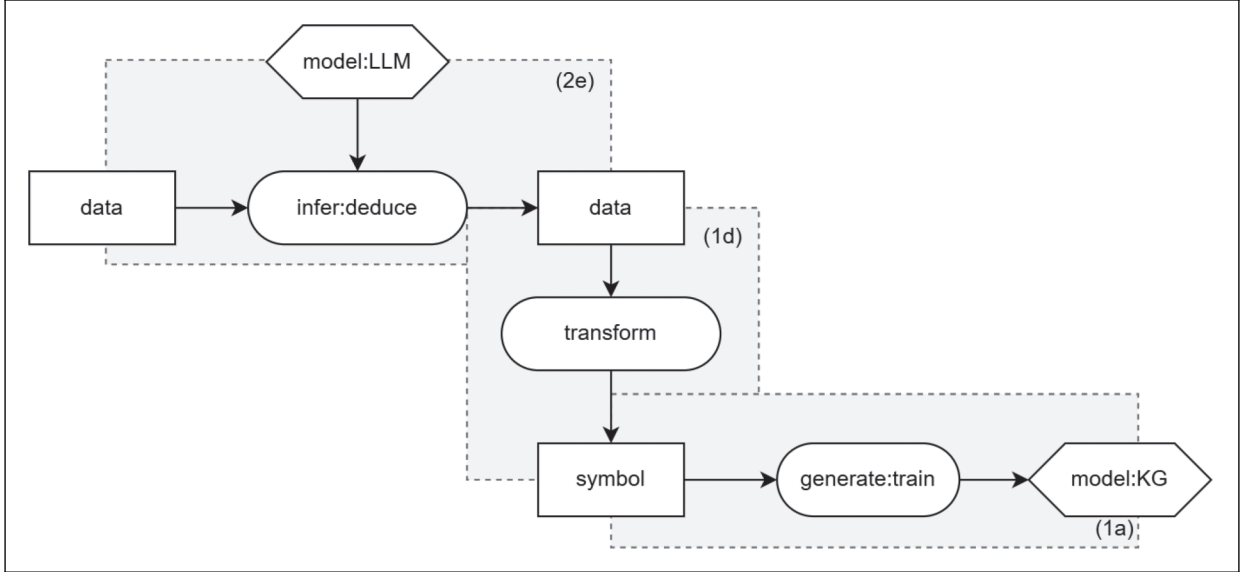


Figure 4. LLM-augmented KGs in training. KG: knowledge graph; LLM: large language model.

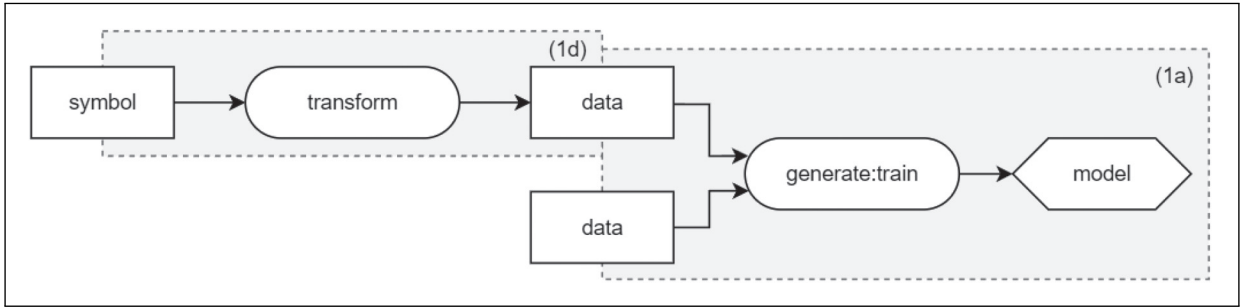


Figure 5. Synergised LLMs and KGs in training. KG: knowledge graph; LLM: large language model.

and pre-trained language representation (Ke et al., 2021; Wang et al., 2021; Yasunaga et al., 2022). They have pre-training tasks to come to a joint knowledge embedding and language modelling optimisation. ERNIE proposes a dual encoder system, consisting of a textual encoder that is fused with the KG encoder (Zhang et al., 2019). BERT-MK has a similar dual encoder, but adds additional information from neighbouring entities to the KG (He et al., 2020). Coke-BERT further improves on this idea by adding a module to filter out irrelevant neighbouring entities (Su et al., 2021). JAKET fuses the entity representation in the middle layers of the LLM (Yu et al., 2022).

4.2 LLM-based Neuro-Symbolic Design Patterns During Inference

Neuro-symbolic systems often combine KGs and LLMs during inference, after training. In this way, the system is more robust to new situations. Many of the LLM-based neuro-symbolic systems follow one of the pre-defined patterns. This section will highlight three depictions of the LLM-based NeSy systems during inference.

4.2.1 KG-enhanced LLMs. KGs can be used to enhance LLMs by utilising the knowledge in KGs. One way to do this is represented in Figure 6. It shows how a KG is used to infer symbols (pattern 2b). These symbols are then transformed to data (pattern 1d) which is used by the LLM to generate new data (pattern 2e). This can be useful for example to align the input data with the knowledge or augment it by adding relevant facts for the LLM to improve the output. In contrast to KG injection during training (see Section 4.1.1), the results of pattern 2b and 1d are now input to the infer process instead of the train process. This means that the knowledge is up to date at the time of inference, rather than at the time of training, which may happen a long time before deployment.

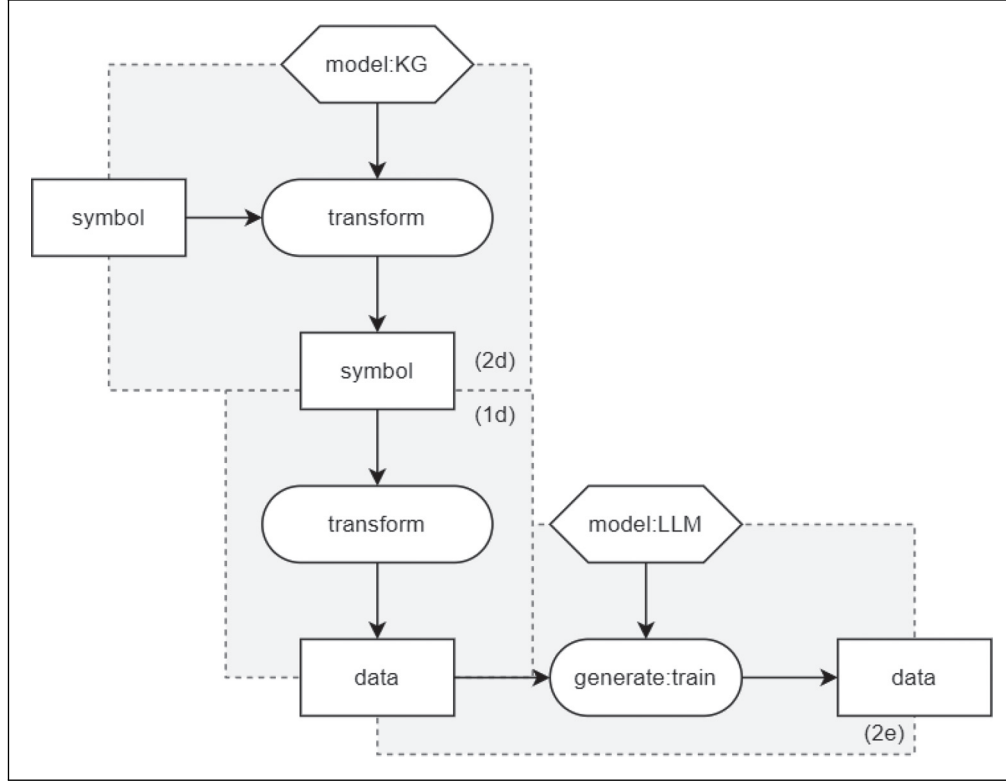


Figure 6. KG-enhanced LLMs during inference. KG: knowledge graph; LLM: large language model.

This pattern describes systems that transform the input data by aligning them with the knowledge of the KG before they are fed into the deduction process with an LLM model. This can be done in a prompt engineering process using KGs (Li et al., 2023; Luo et al., 2023; Wang et al., 2023a; Wen et al., 2023) or retrieval-augmented knowledge methods such as RAG (Lewis et al., 2020). KagNet first encodes the input KG and then augments it with textual representation (Lin et al., 2019).

4.2.2 LLM-Augmented KGs. LLMs can be used to augment KGs to improve information deductions (Figure 7). An LLM is used to infer data (pattern 2e). Then the data is transformed to symbols (pattern 1d) for the KG to reason over (pattern 2b). As with KG-enhanced LLMs during inference, the difference between training and inference for LLM-augmented KGs is that the first pattern is input to infer process rather than to train process of the KG.

One example is using LLMs for KG embedding. Pretrain-KGE uses an LLM to encode the text of the parts of the triples and uses that encoding as a starting point for the KG encoding (Zhang et al., 2020). Moreover, in answering questions with LLM-augmented KG, LLMs are used to bridge the gap between natural language questions and the retrieval of KG answers (Hu et al., 2023; Lukovnikov et al., 2019). In addition, LLMs can be used for the generation of text from a KG, where LLMs are used to generate natural language that describes facts from KGs (Feng et al., 2023; Sun et al., 2023; Wang et al., 2024b). MHGRN uses the LLM representation of the text to guide the reasoning process in the KGs (Feng et al., 2020).

4.2.3 Synergised LLMs and KGs. LLMs and KGs can be combined to work in synergy, also in the application phase. Figure 8 shows how this can be applied, specifically in the case of synergised reasoning. Here, the model is fed both symbols and data, both in the training and in the application phase.

Examples of such methods are JointLK (Sun et al., 2022) and GreaseLM (Zhang et al., 2022). They include interactions between the tokens in the textual input and the entities in the graph in the model's layers. QA-GNN (Yasunaga et al., 2021) represents the LLM information as a special node in the KG for reasoning.

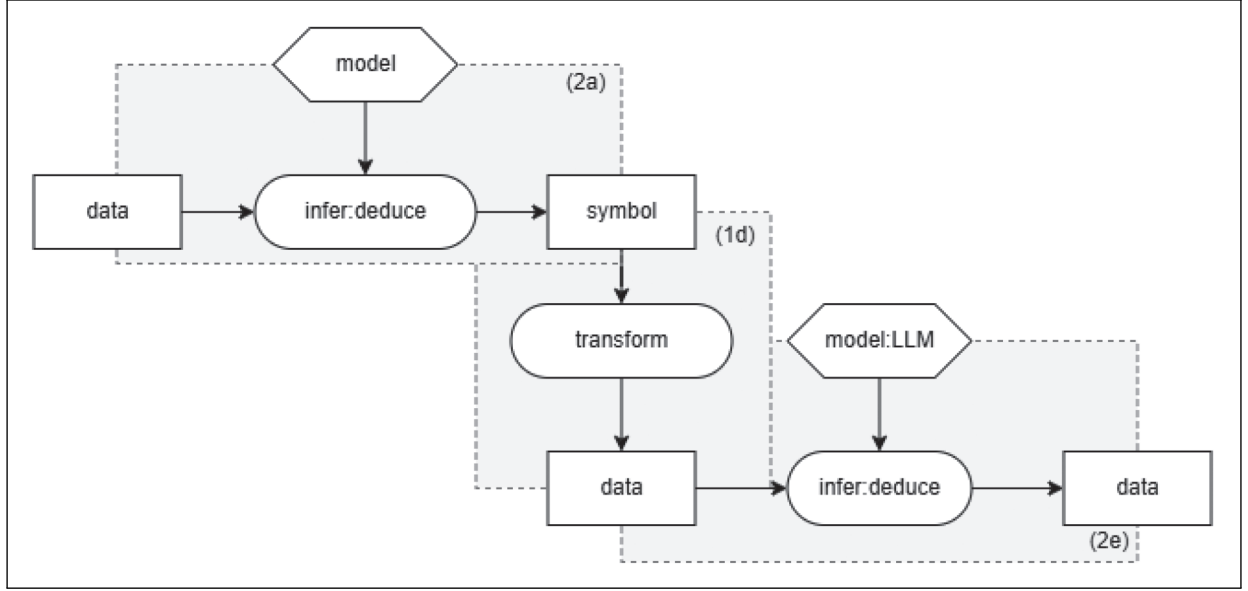


Figure 9. Use of retrieval-augmented generation.

5.1 Retrieval-augmented generation

RAG is a method that expands an LLM with external knowledge (Lewis et al., 2020). A RAG system has two main components, a retriever and a generator. Figure 9 shows the Boxology representation of a RAG system, where the retriever is the model in pattern 2a and 1d and the generator is the LLM in pattern 2e. Firstly, the retriever selects relevant documents based on the posed question (pattern 2a), through classification or with help of a KG. Secondly, the question and retrieved documents are transformed (1d) to be presented to an LLM in a prompt (pattern 2e). Thirdly, the LLM generates an answer to the question based on the information in the selected documents. The LLM can also present the source of the information, making it more trustworthy and reliable.

In KD-CoT, KSL and Think-on-graph, facts are retrieved from a KG together with the reasoning, and an LLM generates a natural language answer to be presented to the user (Feng et al., 2023; Sun et al., 2023; Wang et al., 2024b). RAG is a prime example of the KG-enhanced LLM pattern, as presented in Section 4.2.1.

5.2 KnowGL

The KnowGL parser is a system developed by IBM Research for converting data into symbols. More specifically, it can be used to automatically extract KGs from collections of text documents (Rossiello et al., 2023). KnowGL employs an LLM to extract semantic triples from each sentence, which are then enriched with semantic annotations. Figure 10 shows the Boxology representation of the KnowGL parser. Pattern 2e represents the BART-large model receiving a sentence and inducing a list ‘subject, relation, object’. In the next step, represented by pattern 2b, a ranked list is created of distinct facts and their scores. In the final step, the generated facts are linked to Wikidata. This is done using a mapping of labels to Wikidata IDs (pattern 2b). In the case that the LLM has created a new entity, type, or relation label that are not in Wikidata it returns ‘null’.

The architecture of the KnowGL parser displays a variation of the LLM-augmented KG in inference pattern, in Section 4.2.2.

5.3 KnowBERT

Although knowledge is mostly injected into statistical generative models during the input or during the output stage, approaches to inject knowledge inside the model have also been proposed. A prominent example is KnowBERT, a modified variant of BERT (Peters et al., 2019). It stands out for its fusion of contextual and graph representations, attention-enhanced entity-spanned knowledge infusion, and flexibility in injecting multiple KGs at various model levels. KnowBERT embeds multiple knowledge bases (KB—WordNet and a subset of Wikipedia) into LLMs to enhance their representations with structured, human-curated knowledge. By integrating the Knowledge Attention and Recontextualisation layers (Balažević

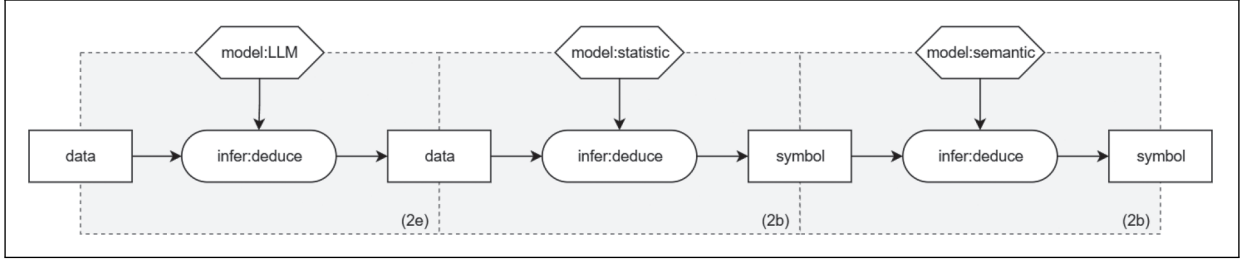


Figure 10. Boxology representation of KnowGL (Rossiello et al., 2023).

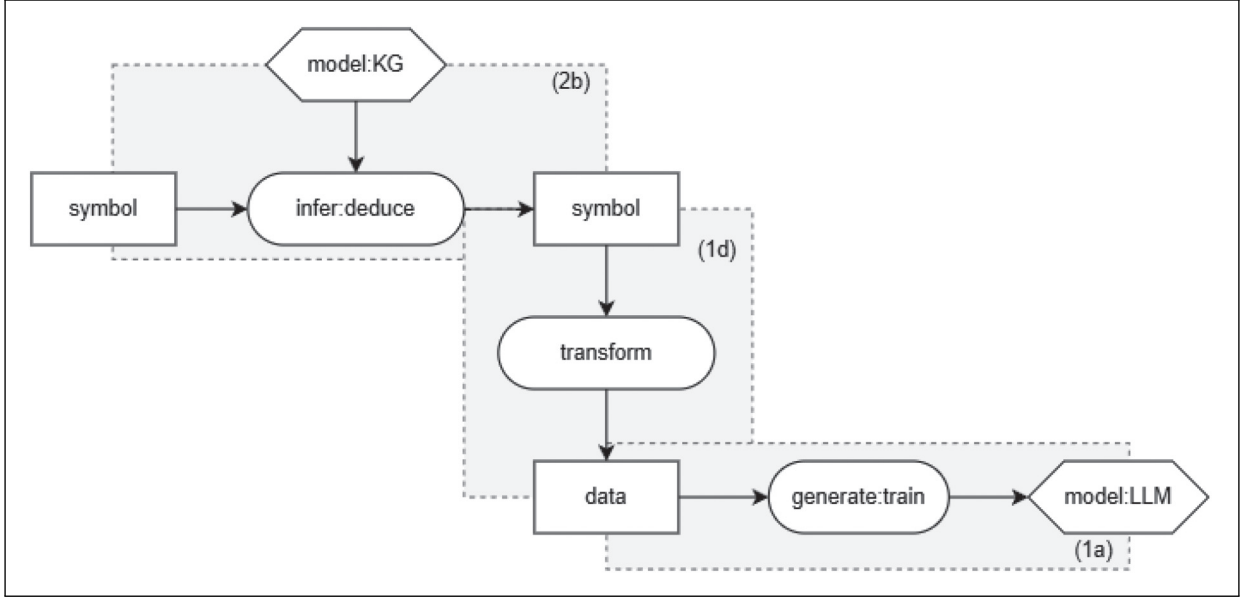


Figure 11. Boxology representation of KnowBERT (Peters et al., 2019).

et al., 2019), graph entity embeddings are used that are processed through an attention mechanism to enhance entity span embeddings. This happens in later layers of the model to stabilise training, but can also potentially be used to inject knowledge at earlier stages (Colon-Hernandez et al., 2021). The Boxology pattern for KnowBERT is shown in Figure 11 and is the same as the pattern presented in Section 4.1.1 (KG-enhanced LLMs in training). Pattern 2b represents the incorporation of KB into a pre-trained BERT model, using an integrated entity linker, as shown with pattern 1d. Finally, the Knowledge Attention and Recontextualisation component is the heart of KnowBert, which is represented as pattern 1a.

5.4 Mathematical Conjecturing and LLMs

Theory explanation or automated conjecturing is the process of inventing new conjectures about a set of functions. The system Johansson and Smallbone (2023) has two principal components: (1) the system assigns the generative task of discovering mathematical conjectures to an LLM, (2) the results are checked using a symbolic theorem prover or counterexample finder. The LLM is first trained on data from a formal language (pattern 3c). The system is then prompted with a formal theory (e.g. a sort function), and has the LLM generate lemmas from the theory as output data. These generated lemmas are transformed from data to symbols (pattern 1d) and are subsequently used by a semantic model(s) prover (pattern 2b). The Boxology representation is depicted in Figure 12.

The approach taken in Yang et al. (2023) is also captured by this representation. The system proposed first uses an LLM component that has been trained on Prolog to generate Prolog code as output (pattern 3c). The output is then transformed to symbols (1d), a symbolic inference engine then produces answers and reasoning traces by executing the code mentioned above (pattern 2b).

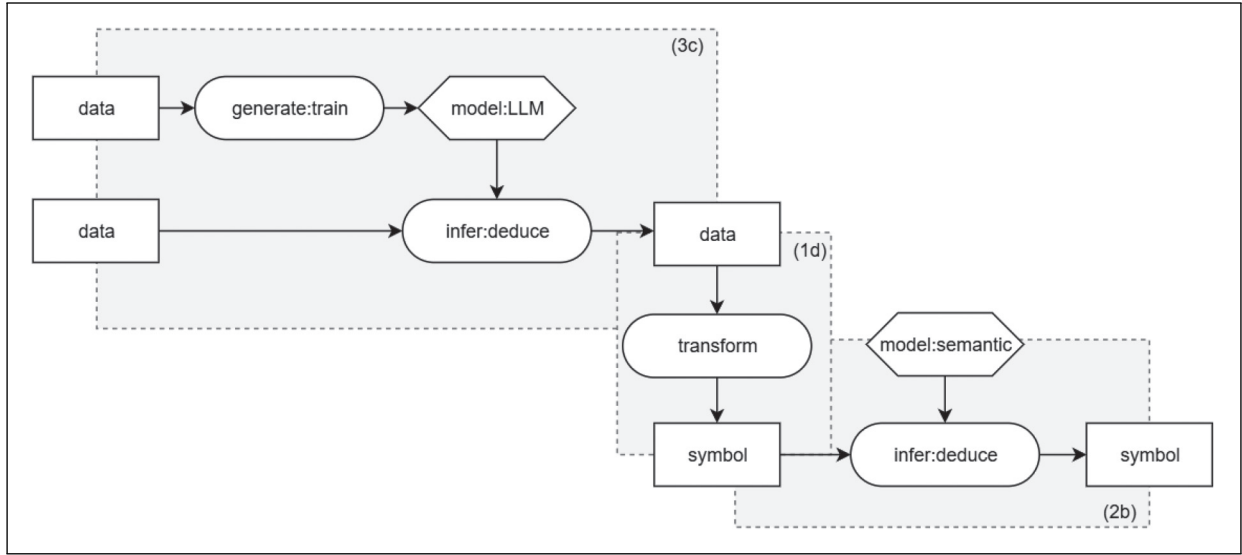


Figure 12. Boxology representation for using large language models (LLMs) for discovery of mathematical conjectures (Johansson & Smallbone, 2023).

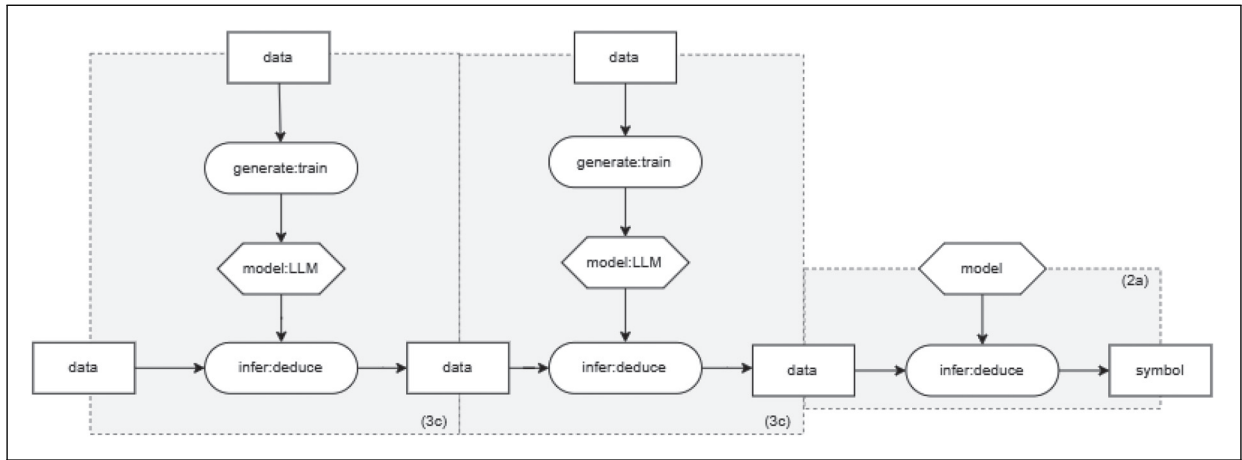


Figure 13. Boxology representation of GENOME (Chen et al., 2023).

Both of these examples show a generalisation of the Boxology pattern in Section 4.2.2 (LLM-augmented KGs during inference), where the KG is replaced by a different semantic model.

5.5 GENOME

Generative Neuro-Symbolic Visual Reasoning by Growing and Reusing Modules (GENOME) focuses on the task of generative software module learning (Chen et al., 2023). Its architecture is based on one LLM generating signatures (input/output) for these software modules and reasoning steps, while another LLM subsequently creates the targeted software module based on those; as usually, both LLMs have been created by means of pre-training. Finally, GENOME employs a deductive reasoner to evaluate the LLM-generated module on a test case.

Figure 13 shows the Boxology representation of GENOME. The system consists of three stages, module initialisation, module generation and module execution, represented by two compositional and one elementary pattern. First, an LLM assesses a visual-language question and outputs new module signatures and operation steps as a response to the query (pattern 3c), if current modules from a library cannot provide an adequate response. In the next step, the LLM generates a module (software code) based on the signature/test case (pattern 3c, 2nd component). Finally, the module is executed by passing it a visual query (pattern 2a).

This use case is an extension of the LLM-augmented KGs during inference as described in Section 4.2.2, deploying two LLMs in sequence.

5.6 Logic-LM

The approach taken by Logic-LM Pan et al. (2023) integrates LLMs as a natural language interface with symbolic solvers to improve logical problem-solving. The logical problem can e.g. be stated as logical programming, first-order logic or a constraint satisfaction problem. This approach is depicted in Boxology notation in Figure 14. Here the system uses LLMs that are trained in the specific logic language (pattern 3c) to translate a problem stated in natural language into a symbolic formulation (patterns 3c and 1d). In the next step, a symbolic reasoner module performs logical inference on the formulated problem and transforms the symbolic results into data (patterns 2b and 1d), using a semantic model and a transformation of logical symbols to a prompt. Finally, an LLM receives the results as an input prompt and outputs a solution in natural language (pattern 3c). The LLM thus functions as a fluent language interface to and from a symbolic reasoner component. The main reasoning is performed by a logic engine (symbolic reasoner), in order to guarantee correct and verifiable results.

This representation is a combination of the patterns presented in LLM-augmented KGs during inference, Section 4.2.2, and KG-enhanced LLMs during inference, Section 4.2.1.

6 Discussion

Advancements. A key contribution of this work is the introduction of novel design patterns. In particular, we introduce a novel elementary pattern and a novel compositional pattern to complement the existing Boxology patterns. The novel patterns proposed in this paper provide a more fine-grained description of the model element, introducing the generative model. The novel elementary pattern allows, for the first time, to use the Boxology to represent data in - data out processing, which is the key concept underlying generative AI techniques. Although we apply it exclusively to LLMs, we have argued that its level of abstraction lends itself to representing other types of generative models, such as generative image models. As a natural derivative of this, the newly introduced compositional pattern employed this elementary pattern to integrate training and inference for generative models. Providing three principled integrations of LLMs and KGs, we illustrate how the composition of elementary patterns can be used to describe LLMs, and we explore several categories as well as specific approaches in use cases, such as KnowGL, GENOME and Logic-LM.

Conceptualisation. Whereas our extension demonstrates its applicability in describing key architectural features of LLM-based neuro-symbolic systems, it raises interesting questions about the nature of other Boxology elements directly related to the new element representing generative models (`model:generative`). The classical opposition of data and symbols, for example, has been a long-standing premise of the Boxology framework. Considering the input and output related to generative models, however, one may wonder whether the output is actually of the same nature. While no straightforward answer seems to be at hand, the framework might benefit from revisiting the underlying nature of data versus symbols. At the same time, we notice that the deductive inference element perhaps generalises from the finer details of the inference capabilities of a generative model to a certain extent. The probabilistic nature of these models may well be captured by a finer distinction in the Boxology inference concept, which can perhaps also cover new classes of inference that have so far not been covered.

Ambiguity. In the process of applying basic design patterns to specific use cases, naturally questions arise about which pattern combinations are allowed and which are not. However, the truth is that Boxology will not in all real-world LLM-based neuro-symbolic systems lead to one unique and unambiguous representation. In practical classroom settings, for example, we have experienced the updated Boxology framework to yield multiple options for describing the same system, depending on the individual student or engineer. Much like some of the notations of unified modelling language (UML) have proven to lend itself to different designs of the same system based on different points of view, this phenomenon is likewise inherent to the Boxology framework. In order to reduce some of the ambiguity, a more formal description along with guidelines for interpretation and use of the current set of Boxology elements seems warranted.

7 Conclusion and Future Work

Despite many open questions and challenges towards generative AI techniques, LLMs are widely used in a variety of applications nowadays. To this end, combining data-driven approaches with knowledge-based techniques is a promising development to address these challenges. In this paper, we propose new design patterns for modular LLM-based neuro-symbolic systems to be included in the design pattern approach for neuro-symbolic systems as proposed by van Bekkum

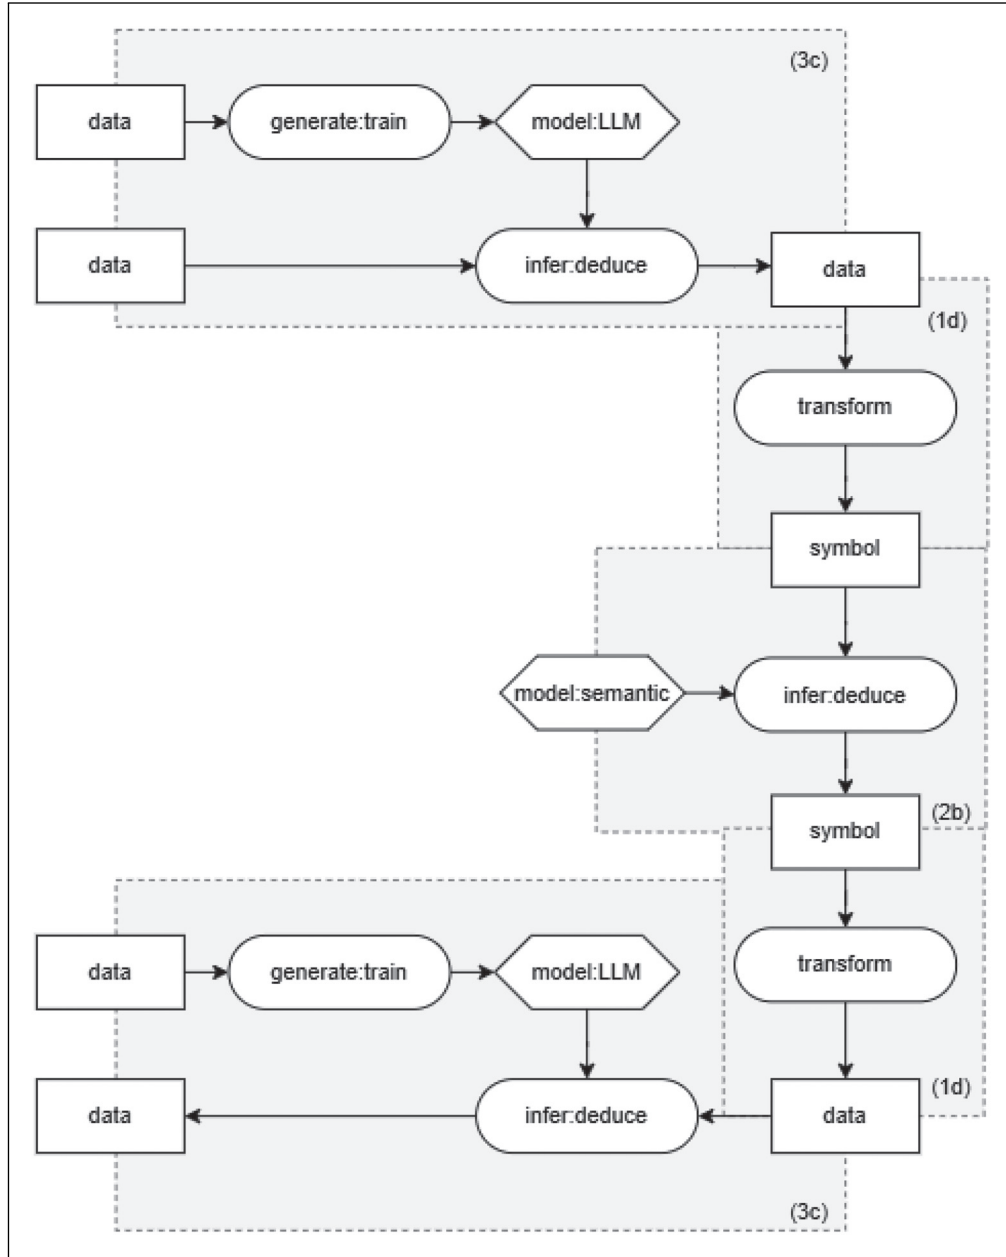


Figure 14. Boxology representation of Logic-LM (Pan et al., 2023).

et al. (2021). Thereby, we are filling a gap in the classic Boxology framework with respect to the recent rise of generative AI techniques. To this end, this paper proposes an extension in terms of concepts and patterns to the set of Boxology elements and patterns as described in earlier work on Boxology (de Boer et al., 2024; van Bekkum et al., 2021; Van Harmelen & Ten Teije, 2019). Specifically, our goal is to make the Boxology framework compatible with the concepts underlying LLMs and LLM-based neuro-symbolic systems. Given the fact that many existing, as well as many potential real-world applications are based on this disruptive AI paradigm, the extension provided within this paper can be considered a substantial conceptual update allowing to maintain the Boxology framework's relevance in the years to come.

In future work, we plan to further explore the usability and benefits of the updated Boxology and its design patterns in domains adjacent to LLM-based neuro-symbolic systems, such as generative AI systems without symbolic AI attached and multi-modal generative AI systems. Moreover, we anticipate the need to further extend and deepen the Boxology framework itself. Temporal or recurring/iterative aspects, for example, have not yet been taken into account and can currently not be visualised adequately in this respect. Our current investigation has also shown that the current concept naming, concept

labelling and some of the formalisation of the Boxology could benefit from critical review and in-depth revisiting. The importance of representing or modelling datasets, for example, may be taken into account in future specifications of particular subtypes of instances and models. Finally, we consider a future use of graphical tools for the Boxology beneficial. In software development, this approach has proven both efficient and effective and is well known, for example, from the UML and visual programming tools, such as LabView or Scratch. And while our current work is mostly concerned with graphical representations of design patterns for system design and documentation, the promise of templates, low-code, or no-code development, seems also an appealing field of research for the future.

Acknowledgements

We thank Frank van Harmelen and Annette ten Teije for their feedback. We also thank Daan Di Scala for his contribution to the KnowGL pattern.


Funding


The authors received the following financial support for the research, authorship, and/or publication of this article: This study was supported by the TNO project GRAIL.


Declaration of conflicting interests


The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

ORCID iDs

Maaïke de Boer  <https://orcid.org/0000-0002-2775-8351>

Quirine Smit  <https://orcid.org/0009-0008-8890-4261>

Michael van Bakkum  <https://orcid.org/0009-0007-3009-254X>

André Meyer-Vitali  <https://orcid.org/0000-0002-5242-1443>

References

- Agrawal, G., Kumarage, T., Alghamdi, Z., & Liu, H. (2023). Can knowledge graphs reduce hallucinations in llms? A survey, *arXiv preprint arXiv:2311.07914*.
- Amador-Domínguez, E., Serrano, E., & Manrique, D. (2024). Neurosymbolic system profiling: A template-based approach. *Knowledge-Based Systems*, 287, 111441. <https://doi.org/10.1016/j.knosys.2024.111441>. <https://www.sciencedirect.com/science/article/pii/S0950705124000765>
- Balažević, I., Allen, C., & Hospedales, T. M. (2019). Tucker: Tensor factorization for knowledge graph completion, *arXiv:1901.09590*.
- Beck, M., Pöppel, K., Spanring, M., Auer, A., Prudnikova, O., Kopp, M., Klambauer, G., Brandstetter, J., & Hochreiter, S. (2024). xLSTM: Extended long short-term memory, *arXiv preprint arXiv:2405.04517*.
- 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., Pinkas, G., & Poon, H. (2021). Neural-symbolic learning and reasoning: A survey and interpretation 1. In *Neuro-symbolic artificial intelligence: The state of the art* (pp. 1–51). IOS Press.
- Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., & Manassra, W. (2023). Improving image generation with better captions. *Computer Science*, 2(3), 8.
- Breit, A., Waltersdorfer, L., Ekaputra, F. J., Sabou, M., Ekelhart, A., Iana, A., Paulheim, H., Portisch, J., Revenko, A., Teije, A. T., & Van Harmelen, F. (2023). Combining machine learning and semantic web: A systematic mapping study. *ACM Computing Surveys*, 55(14s), 1–41.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., & Agarwal, S. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Cattan, A., Eirew, A., Stanovsky, G., Joshi, M., & Dagan, I. (2021). Cross-document coreference resolution over predicted mentions. In C. Zong, F. Xia, W. Li & R. Navigli (Eds.), *Findings of the association for computational linguistics: ACL-IJCNLP 2021* (pp. 5100–5107). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-acl.453>. <https://aclanthology.org/2021.findings-acl.453>
- Chen, H., Jiao, F., Li, X., Qin, C., Ravaut, M., Zhao, R., Xiong, C., & Joty, S. (2023). ChatGPT's one-year anniversary: Are open-source large language models catching up? *arXiv:2311.16989*.
- Chen, Z., Sun, R., Liu, W., Hong, Y., & Gan, C. (2023). Genome: Generative neuro-symbolic visual reasoning by growing and reusing modules, *arXiv preprint arXiv:2311.04901*.
- Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., & Schuh, P. (2023). Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240), 1–113.

- Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, Y., Wang, X., Dehghani, M., Brahma, S., & Webson, A. (2022). Scaling instruction-finetuned language models, *arXiv:2210.11416*.
- Colclough, B. C., & 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)*. Co-located with IJCAI 2024, Jeju, South Korea.
- Colon-Hernandez, P., Havasi, C., Alonso, J., Huggins, M., & Breazeal, C. (2021). Combining pre-trained language models and structured knowledge, *arXiv preprint arXiv:2101.12294*.
- Cordeschi, R. (2007). AI turns fifty: Revisiting its origins. *Applied Artificial Intelligence*, 21(4–5), 259–279.
- de Boer, M., Smit, Q., van Bekkum, M., Meyer-Vitali, A., & Schmid, T. (2024). Modular design patterns for generative neuro-symbolic systems, *GeNeSy*.
- De Raedt, L., Dumančić, S., Manhaeve, R., & Marra, G. (2020). From statistical relational to neuro-symbolic artificial intelligence, *arXiv preprint arXiv:2003.08316*.
- Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805*.
- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., & Goyal, A. (2024). The llama 3 herd of models, *arXiv preprint arXiv:2407.21783*.
- Ellis, A., Dave, B., Salehi, H., Ganapathy, S., & Shimizu, C. (2024a). EASY-AI: sEmantic And compoSable gLYphs for representing AI systems. In *HAI 2024: Hybrid human AI systems for the social good* (pp. 105–113). IOS Press.
- Ellis, A., Dave, B., Salehi, H., Ganapathy, S., & Shimizu, C. (2024b). Implementing SNOOP-AI in CoModIDE. In *NAECON 2024-IEEE National Aerospace and Electronics conference* (pp. 101–104). IEEE.
- Feng, C., Zhang, X., & Fei, Z. (2023). Knowledge solver: Teaching llms to search for domain knowledge from knowledge graphs, *arXiv preprint arXiv:2309.03118*.
- Feng, Y., Chen, X., Lin, B. Y., Wang, P., Yan, J., & Ren, X. (2020). Scalable multi-hop relational reasoning for knowledge-aware question answering. In *Proceedings of EMNLP* (pp. 1295–1309). Association for Computational Linguistics.
- Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey, *arXiv preprint arXiv:2312.10997*.
- Garcez, A.D., & Lamb, L. C. (2023). Neurosymbolic AI: The 3 rd wave. *Artificial Intelligence Review*, 56(11), 12387–12406.
- He, B., Zhou, D., Xiao, J., Jiang, X., Liu, Q., Yuan, N. J., & Xu, T. (2020). BERT-MK: Integrating graph contextualized knowledge into pre-trained language models. In T. Cohn, Y. He & Y. Liu (Eds.), *Findings of the association for computational linguistics: EMNLP 2020* (pp. 2281–2290). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.207>. <https://aclanthology.org/2020.findings-emnlp.207>
- Hilario, M. (1994). An overview of strategies for neurosymbolic integration. In R. Sun & L. Bookman (Eds.), *Computational architectures integrating symbolic and neural processes*. Kluwer Academic Publishers.
- Hitzler, P., Eberhart, A., Ebrahimi, M., Sarker, M. K., & Zhou, L. (2022). Neuro-symbolic approaches in artificial intelligence. *National Science Review*, 9(6), nwac035. <https://doi.org/10.1093/nsr/nwac035>
- Hu, N., Wu, Y., Qi, G., Min, D., Chen, J., Pan, J. Z., & Ali, Z. (2023). An empirical study of pre-trained language models in simple knowledge graph question answering. *World Wide Web*, 26(5), 2855–2886. <https://doi.org/10.1007/s11280-023-01166-y>
- Huang, N., Deshpande, Y. R., Liu, Y., Alberts, H., Cho, K., Vania, C., & Calixto, I. (2022). Endowing language models with multimodal knowledge graph representations, *arXiv preprint arXiv:2206.13163*.
- Huang, X., Ruan, W., Huang, W., Jin, G., Dong, Y., Wu, C., Bensalem, S., Mu, R., Qi, Y., Zhao, X., Cai, K., Zhang, Y., Wu, S., Xu, P., Wu, D., Freitas, A., & Mustafa, M. A. (2024). A survey of safety and trustworthiness of large language models through the lens of verification and validation. *Artificial Intelligence Review*, 57(7), 175. <https://doi.org/10.1007/s10462-024-10824-0>
- Ji, S., Pan, S., Cambria, E., Marttinen, P., & Philip, S. Y. (2021). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2), 494–514.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 1–38.
- Jin, Y., Li, J., Liu, Y., Gu, T., Wu, K., Jiang, Z., He, M., Zhao, B., Tan, X., Gan, Z., & Wang, Y. (2024). Efficient multimodal large language models: A survey, *arXiv preprint arXiv:2405.10739*.
- Johansson, M., & Smallbone, N. (2023). Exploring mathematical conjecturing with large language models. In *Proceedings of NeSy*.
- Joshi, M., Chen, D., Liu, Y., Weld, D. S., Zettlemoyer, L., & Levy, O. (2020). Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8, 64–77.
- Kaddour, J., Harris, J., Mozes, M., Bradley, H., Raileanu, R., & McHardy, R. (2023). Challenges and applications of large language models, *arXiv preprint arXiv:2307.10169*.
- Ke, P., Ji, H., Ran, Y., Cui, X., Wang, L., Song, L., Zhu, X., & Huang, M. (2021). JointGT: Graph-text joint representation learning for text generation from knowledge graphs. In C. Zong, F. Xia, W. Li & R. Navigli (Eds.), *Findings of the association for computational linguistics: ACL-IJCNLP 2021* (pp. 2526–2538). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-acl.223>. <https://aclanthology.org/2021.findings-acl.223>

- Khanday, N. Y., & Sofi, S. A. (2021). Taxonomy, state-of-the-art, challenges and applications of visual understanding: A review. *Computer Science Review*, 40, 100374. <https://doi.org/https://doi.org/10.1016/j.cosrev.2021.100374>. <https://www.sciencedirect.com/science/article/pii/S1574013721000149>
- Kim, B., Hong, T., Ko, Y., & Seo, J. (2020). Multi-task learning for knowledge graph completion with pre-trained language models. In D. Scott, N. Bel & C. Zong (Eds.), *Proceedings of the 28th international conference on computational linguistics* (pp. 1737–1743). International Committee on Computational Linguistics. <https://doi.org/10.18653/v1/2020.coling-main.153>. <https://aclanthology.org/2020.coling-main.153>
- Kukreja, S., Kumar, T., Purohit, A., Dasgupta, A., & Guha, D. (2024). A literature survey on open source large language models. In *Proceedings of the 2024 7th international conference on computers in management and business* (pp. 133–143). Association for Computing Machinery. ISBN 9798400716652.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., & Riedel, S. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- Li, S., Gao, Y., Jiang, H., Yin, Q., Li, Z., Yan, X., Zhang, C., & Yin, B. (2023). Graph reasoning for question answering with triplet retrieval, *arXiv preprint arXiv:2305.18742*.
- Li, S., Li, X., Shang, L., Sun, C., Liu, B., Ji, Z., Jiang, X., & Liu, Q. (2022). Pre-training language models with deterministic factual knowledge, *arXiv preprint arXiv:2210.11165*.
- Lin, B. Y., Chen, X., Chen, J., & Ren, X. (2019). KagNet: Knowledge-aware graph networks for commonsense reasoning. In K. Inui, J. Jiang, V. Ng & X. Wan (Eds.), *Proceedings of EMNLP-IJCNLP* (pp. 2829–2839). Association for Computational Linguistics.
- Lin, Z., Guan, S., Zhang, W., Zhang, H., Li, Y., & Zhang, H. (2024). Towards trustworthy LLMs: A review on debiasing and dehallucinating in large language models. *Artificial Intelligence Review*, 57(9), 243.
- Lukovnikov, D., Fischer, A., & Lehmann, J. (2019). Pretrained transformers for simple question answering over knowledge graphs. In *The semantic web – ISWC 2019: 18th international semantic web conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part I* (pp. 470–486). ISBN 978-3-030-30792-9.
- Luo, L., Ju, J., Xiong, B., Li, Y.-F., Haffari, G., & Pan, S. (2023). Chatrule: Mining logical rules with large language models for knowledge graph reasoning, *arXiv preprint arXiv:2309.01538*.
- Lv, X., Lin, Y., Cao, Y., Hou, L., Li, J., Liu, Z., Li, P., & Zhou, J. (2022). Do pre-trained models benefit knowledge graph completion? A reliable evaluation and a reasonable approach. In S. Muresan, P. Nakov & A. Villavicencio (Eds.), *Findings of the association for computational linguistics: ACL 2022* (pp. 3570–3581). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2022.findings-acl.282>. <https://aclanthology.org/2022.findings-acl.282>
- 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., & 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.
- Meyer-Vitali, A., Mulder, W., & de Boer, M. H. (2021). Modular design patterns for hybrid actors, *arXiv preprint arXiv:2109.09331*.
- Min, B., Ross, H., Sulem, E., Veyseh, A. P. B., Nguyen, T. H., Sainz, O., Agirre, E., Heintz, I., & Roth, D. (2023). Recent advances in natural language processing via large pre-trained language models: A survey. *ACM Computing Surveys*, 56(2), 1–40.
- Minaee, S., Mikolov, T., Nikzad, N., Chenaghlu, M., Socher, R., Amatriain, X., & Gao, J. (2024). Large language models: A survey, *arXiv preprint arXiv:2402.06196*.
- Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1), 8–30.
- Mossakowski, T. (2022). Modular design patterns for neural-symbolic integration: Refinement and combination, *arXiv:2206.04724*.
- Muennighoff, N., Wang, T., Sutawika, L., Roberts, A., Biderman, S., Scao, T. L., Bari, M. S., Shen, S., Yong, Z.-X., Schoelkopf, H., & Tang, X. (2022). Crosslingual generalization through multitask finetuning, *arXiv:2211.01786*.
- Nayyeri, M., Wang, Z., Akter, M. M., Alam, M. M., Rony, M. R. A. H., Lehmann, J., & Staab, S. (2023). Integrating knowledge graph embeddings and pre-trained language models in hypercomplex spaces. In *International semantic web conference* (pp. 388–407). Springer.
- OECD. (2023). *Is education losing the race with technology? AI's progress in maths and reading*. In *Educational research and innovation*. OECD Publishing. <https://doi.org/10.1787/73105f99-en>
- Pan, L., Albalak, A., Wang, X., & Wang, W. Y. (2023). Logic-lm: Empowering large language models with symbolic solvers for faithful logical reasoning, *arXiv:2305.12295*.
- Pan, S., Luo, L., Wang, Y., Chen, C., Wang, J., & Wu, X. (2024). Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*, 36(7), 3580–3599.
- Park, S., & Kim, H. (2021). Improving sentence-level relation extraction through curriculum learning, *arXiv preprint arXiv:2107.09332*.

- Peng, B., Alcaide, E., Anthony, Q., Albalak, A., Arcadinho, S., Biderman, S., Cao, H., Cheng, X., Chung, M., Grella, M., & Gv, K. K. (2023). RWKV: Reinventing RNNs for the transformer era, *arXiv preprint arXiv:2305.13048*.
- Peters, M. E., Neumann, M., Logan IV, R. L., Schwartz, R., Joshi, V., Singh, S., & Smith, N. A. (2019). Knowledge enhanced contextual word representations, *arXiv:1909.04164*.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8), 9.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2021). High-resolution image synthesis with latent diffusion models.
- Rosset, C., Xiong, C., Phan, M., Song, X., Bennett, P., & Tiwary, S. (2020). Knowledge-aware language model pretraining, *arXiv preprint arXiv:2007.00655*.
- Rossiello, G., Chowdhury, M. F. M., Mihindukulasooriya, N., Cornec, O., & Gliozzo, A. M. (2023). KnowGL: Knowledge generation and linking from text. In *AAAI* (pp. 16476–16478).
- Sabou, M., Llugiqi, M., Ekaputra, F. J., Waltersdorfer, L., & Tsaneva, S. (2024). Knowledge engineering in the age of neurosymbolic systems, *Neurosymbolic AI Journal* (under review).
- Samsi, S., Zhao, D., McDonald, J., Li, B., Michaleas, A., Jones, M., Bergeron, W., Kepner, J., Tiwari, D., & Gadepally, V. (2023). From words to watts: Benchmarking the energy costs of large language model inference. In *2023 IEEE high performance extreme computing conference (HPEC)* (pp. 1–9). IEEE.
- Savage, N. (2020). The race to the top among the world’s leaders in artificial intelligence. *Nature*, 588(7837), S102–S102.
- Schmid, T. (2023). A systematic and efficient approach to the design of modular hybrid AI systems. In *AAAI spring symposium on challenges requiring the combination of machine learning and knowledge engineering (AAAI-MAKE), CEUR workshop proceedings* (Vol. 3433).
- Schmidt, W. J., Rincon-Yanez, D., Kharlamov, E., & Paschke, A. (2024). Scaling scientific knowledge discovery with neuro-symbolic AI and large language models. In *Proceedings of the first international workshop on scaling knowledge graphs for industry, co-located with the 20th international conference on semantic systems (SEMANTICS), Amsterdam, Netherlands*.
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63.
- Shen, T., Mao, Y., He, P., Long, G., Trischler, A., & Chen, W. (2020). Exploiting structured knowledge in text via graph-guided representation learning, *arXiv preprint arXiv:2004.14224*.
- Shi, P., & Lin, J. (2019). Simple bert models for relation extraction and semantic role labeling, *arXiv preprint arXiv:1904.05255*.
- Su, Y., Han, X., Zhang, Z., Lin, Y., Li, P., Liu, Z., Zhou, J., & Sun, M. (2021). CokeBERT: Contextual knowledge selection and embedding towards enhanced pre-trained language models. *AI Open*, 2, 127–134.
- Sun, J., Xu, C., Tang, L., Wang, S., Lin, C., Gong, Y., Shum, H.-Y., & Guo, J. (2023). Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph.
- Sun, R., & Alexandre, F. (1997) *Connectionist-symbolic integration: From unified to hybrid approaches*, 1st Edn. New York: Psychology Press.
- Sun, Y., Shi, Q., Qi, L., & Zhang, Y. (2022). JointLK: Joint reasoning with language models and knowledge graphs for commonsense question answering (pp. 5049–5060).
- Tian, H., Gao, C., Xiao, X., Liu, H., He, B., Wu, H., Wang, H., & Wu, F. (2020). SKEP: Sentiment knowledge enhanced pre-training for sentiment analysis, *arXiv preprint arXiv:2005.05635*.
- van Bekkum, M., de Boer, M., van Harmelen, F., Meyer-Vitali, A., & 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 P. Hitzler & M.K. Sarker (Eds.), *Neuro-symbolic artificial intelligence: The state of the art*. IOS Press, p. V–VII. <https://doi.org/10.3233/FAIA210347-fm>
- Van Harmelen, F., & Ten Teije, A. (2019). A boxology of design patterns for hybrid learning and reasoning systems. *Journal of Web Engineering*, 18(1–3), 97–123.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 5998.
- Wang, J., Jiang, H., Liu, Y., Ma, C., Zhang, X., Pan, Y., Liu, M., Gu, P., Xia, S., Li, W., & Zhang, Y. (2024a). A comprehensive review of multimodal large language models: Performance and challenges across different tasks, *arXiv preprint arXiv:2408.01319*.
- Wang, J., Sun, Q., Li, X., & Gao, M. (2023a). Boosting language models reasoning with chain-of-knowledge prompting, *arXiv preprint arXiv:2306.06427*.
- Wang, P., Xie, X., Wang, X., & Zhang, N. (2023b). Reasoning through memorization: Nearest neighbor knowledge graph embeddings. In *CCF international conference on natural language processing and Chinese computing* (pp. 111–122). Springer.
- Wang, X., Gao, T., Zhu, Z., Zhang, Z., Liu, Z., Li, J., & Tang, J. (2021). KEPLER: A unified model for knowledge embedding and pre-trained language representation. *Transactions of the Association for Computational Linguistics*, 9, 176–194.
- Wang, X., He, Q., Liang, J., & Xiao, Y. (2022). Language models as knowledge embeddings, *arXiv preprint arXiv:2206.12617*.

- Wang, Y., Lipka, N., Rossi, R., Siu, A., Zhang, R., & Derr, T. (2024b). Knowledge graph prompting for multi-document question answering. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38, 19206–19214.
- Wei, X., Wang, S., Zhang, D., Bhatia, P., & Arnold, A. (2021). Knowledge enhanced pretrained language models: A comprehensive survey, *arXiv:2110.08455*.
- Wen, Y., Wang, Z., & Sun, J. (2023). Mindmap: Knowledge graph prompting sparks graph of thoughts in large language models, *arXiv preprint arXiv:2308.09729*.
- Witschel, H. F., Barroca, B., Correia, A., Martin, A., & Caldera, C. (2020). Visualization of patterns for hybrid learning and reasoning with human involvement. In *Proceedings of the 12th international conference on agents and artificial intelligence* (pp. 423–430). SCITEPRESS. <https://doi.org/10.5220/0008975504230430>
- Xie, X., Li, Z., Wang, X., Xi, Z., & Zhang, N. (2022). Lambdakg: A library for pre-trained language model-based knowledge graph embeddings, *arXiv preprint arXiv:2210.00305*.
- Xiong, W., Du, J., Wang, W. Y., & Stoyanov, V. (2019). Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model, *arXiv preprint arXiv:1912.09637*.
- Yan, H., Gui, T., Dai, J., Guo, Q., Zhang, Z., & Qiu, X. (2021). A unified generative framework for various NER subtasks. In C. Zong, F. Xia, W. Li & R. Navigli (Eds.), *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (Volume 1: Long Papers)* (pp. 5808–5822). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.acl-long.451>. <https://aclanthology.org/2021.acl-long.451>
- Yang, S., Li, X., Cui, L., Bing, L., & Lam, W. (2023). Neuro-symbolic integration brings causal and reliable reasoning proofs.
- Yao, L., Mao, C., & Luo, Y. (2019). KG-BERT: BERT for knowledge graph completion, *arXiv preprint arXiv:1909.03193*.
- Yasunaga, M., Bosselut, A., Ren, H., Zhang, X., Manning, C. D., Liang, P. S., & Leskovec, J. (2022). Deep bidirectional language-knowledge graph pretraining. *Advances in Neural Information Processing Systems*, 35, 37309–37323.
- Yasunaga, M., Ren, H., Bosselut, A., Liang, P., & Leskovec, J. (2021). QA-GNN: Reasoning with language models and knowledge graphs for question answering. In *Proceedings of the 2021 conference of the North American chapter of the association for computational linguistics: Human language technologies* (pp. 535–546). Association for Computational Linguistics.
- Yu, D., Yang, B., Liu, D., & Wang, H. (2021). A survey on neural-symbolic learning systems.
- Yu, D., Zhu, C., Yang, Y., & Zeng, M. (2022). Jaket: Joint pre-training of knowledge graph and language understanding. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 36, pp. 11630–11638).
- Zhang, D., Yu, Y., Li, C., Dong, J., Su, D., Chu, C., & Yu, D. (2024). MM-LLMs: Recent advances in multimodal large language models, *arXiv preprint arXiv:2401.13601*.
- Zhang, S., Dong, L., Li, X., Zhang, S., Sun, X., Wang, S., Li, J., Hu, R., Zhang, T., Wu, F., & Guoyin, W. (2023). Instruction tuning for large language models: A survey, *arXiv preprint arXiv:2308.10792*.
- Zhang, X., Bosselut, A., Yasunaga, M., Ren, H., Liang, P., Manning, C. D., & Leskovec, J. (2022). GreaseLM: Graph reasoning enhanced language models for question answering, *arXiv preprint arXiv:2201.08860*.
- Zhang, Z., Han, X., Liu, Z., Jiang, X., Sun, M., & Liu, Q. (2019). ERNIE: Enhanced language representation with informative entities, *arXiv preprint arXiv:1905.07129*.
- Zhang, Z., Liu, X., Zhang, Y., Su, Q., Sun, X., & He, B. (2020). Pretrain-KGE: Learning knowledge representation from pre-trained language models. In T. Cohn, Y. He & Y. Liu (Eds.), *Findings of the association for computational linguistics: EMNLP 2020* (pp. 259–266). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.25>. <https://aclanthology.org/2020.findings-emnlp.25>
- Zhao, H., Chen, H., Yang, F., Liu, N., Deng, H., Cai, H., Wang, S., Yin, D., & Du, M. (2024). Explainability for large language models: A survey. *ACM Transactions on Intelligent Systems and Technology*, 15(2), 1–38.
- Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., & Du, Y. (2023). A survey of large language models, *arXiv preprint arXiv:2303.18223*. <https://arxiv.org/abs/2303.18223>
- Zhong, W., Cui, R., Guo, Y., Liang, Y., Lu, S., Wang, Y., Saied, A., Chen, W., & Duan, N. (2023). Agieval: A human-centric benchmark for evaluating foundation models, *arXiv preprint arXiv:2304.06364*.