

Neuro-Symbolic Knowledge Engineering


Emerging Paradigms and Challenges



Marta Sabou




This talk would not have been possible without: the support of the WU




This talk would not have been possible without: the SemSys TEAM



What we do: SemSys research landscape



Disciplines covered by this talk



Motivation

What is Knowledge Engineering (KE)?
Why should we care about knowledge (again)?

Neuro-symbolic KE

What is neuro-symbolic AI?
How does it affect KE?

Challenges and Emerging Paradigms

What are KE challenges?
What are emerging paradigms from neuro-symbolic KE?

Part 1: Motivation



What is Knowledge Engineering?

“the collection of activities for eliciting, capturing, conceptualising and formalising knowledge for the purpose of being used in information systems”

Expert
Systems

Semantic Web,
Ontologies

Ontology
Networks

Linked Data,
Knowledge Graphs

1990-2000

2001

2012

2022


CommonKADS
(Schreiber et al.)

Ontology 101
(Noy &
McGuinness)

NeOn
Methodology
(Suárez-Figueroa,
et al)

Linked Open Terms
(Poveda-Villalón et al.)


Knowledge Graph (KG) Development Process*: as emerging from practice



Synthesized from 57 papers on KG development

**) Tamašauskaitė, Gytė, and Paul Groth. Defining a knowledge graph development process through a systematic review. ACM Transactions on Software Engineering and Methodology 32(1), 2023.*

Knowledge is a key element of Complex Information Systems




Explaining Events in CPS:

Slow charging at C1:

- * overload at transformer T1
- * caused by houses consuming a high amount of energy
- * because of low energy prices and low batteries of some houses
- * because producing less energy as usual due to reduced solar radiation


Complex Information Systems: Key Characteristics

- Many & heterogenous data sources
- Big data (high volume, high frequency)
- Varying data quality
- Complex domain knowledge @ experts



Diverse set of end users requiring personalization of results

Knowledge is a key element for GenAI/LLMs



S. Pan et. al, Unifying Large Language Models and Knowledge Graphs: A Roadmap. IEEE Trans. on Knowl. and Data Eng. 36(7), 2024.

Knowledge is a key element of Agentic Systems

Knowledge/Context engineering is key to bring AI Agents to production!

- [Deloitte, Nov'25](#): Context layer: This robust **knowledge engineering** foundation is important for **scalable AI agent architecture**. ... well-governed KR (e.g. **KGs, ontologies, domain taxonomies**, ...)
- [McKinsey, Apr'26](#): The **semantic layer** turns data into knowledge ... Without this **shared semantic foundation**, agents may act on incomplete or conflicting interpretations of the same data, increasing error rates and operational risk

Microsoft mentions semantic backbone in the [Fabric IQ announcement](#):

*Fabric IQ combines five integrated capabilities into one semantic intelligence system: **Ontology**, **Semantic Model** (BI), **Graph** (native graph engine), **Data Agent** (virtual analysts that answer business questions), **Operations Agent** (autonomous agents that reason, learn, and act)*

*Fabric IQ ... becomes the **semantic backbone** for every team, application, and AI agent.*


Slide adapted from Atanas Kyriakov, Keynote at ESWC'26

Part 2: Neuro-Symbolic Knowledge Engineering




The Rise of Neuro-symbolic AI

Symbolic AI




- + symbolic knowledge
- + explainability
- vulnerable to noisy data
- knowledge acquisition bottleneck

Subsymbolic AI



- + knowledge from sparse data
- + broad applicability
- low transparency
- lots of training data required

Semantic Web Machine Learning Systems



Our scope, subset/related area to neuro-symbolic AI

→ Dynamic and impactful research area, yet **systematic knowledge missing** in terms of characteristics and system architectures

Review of 476 papers of systems that combine Semantics Web and Machine Learning

Home > ACM Journals > ACM Computing Surveys > Vol. 55, No. 14s > Combining Machine Learning and Semantic Web: A Systematic Mapping Study

SURVEY | OPEN ACCESS |  



Combining Machine Learning and Semantic Web: A Systematic Mapping Study

Authors:  [Anna Breit](#),  [Laura Waltersdorfer](#),  [Fajar J. Ekaputra](#),  [Marta Sabou](#),  [Andreas Ekelhart](#), 
[Andreea Iana](#),  [Heiko Paulheim](#),  [Jan Portisch](#),  [Artem Revenko](#),  [Annette Ten Teije](#),  [Frank Van Harmelen](#) | [Authors Info & Claims](#)

ACM Computing Surveys, Volume 55, Issue 14s • Article No.: 313, Pages 1 - 41 • <https://doi.org/10.1145/3586163>

Published: 17 July 2023 [Publication History](#)



WU
WIRTSCHAFTS
UNIVERSITÄT
WIEN VIENNA
UNIVERSITY OF
ECONOMICS
AND BUSINESS




 **EQUIS**
ACCREDITED  **AACSB**
ACCREDITED  **AMBA**
ACCREDITED

The Rise of Neuro-symbolic AI: AI System Patterns - Examples

A1: s-M-s


**Abstract
Pattern**



A2: d-M-s



**System
(Pattern
Instance)**




Agapito G, Cannataro M, Guzzi PH, et al. Using GO-WAR for mining cross-ontology weighted association rules. *Computer Methods and Programs in Biomedicine*. 2015;120(2):113–122.

Zafar B, Cochez M, Qamar U. Using Distributional Semantics for Automatic Taxonomy Induction. In: *Int. Conf. on Frontiers of Information Technology (FIT)*; 2016.


The Rise of Neuro-symbolic AI: AI System Patterns - Classification

Atomic-Patterns


A1: s-M-s



A2: d-M-s




A3: s-M-d




Fusion-Patterns


F1: d/s-M-d



F2: d/s-M-s




F3: d/s-M-d/s




T-Patterns


T1: {d-M-d/s}-M-s




T2: {d/s}-M-d-M-s



T3: {s-M-d/d}-M-s




T4: {s-M-d/d}-M-d



T5: {d-M-s/s}-K-s




T6: {d-M-d/s}-M-d




I-Patterns

I1: s-M-d-M-s




I2: s-M-d-M-d




Y-Patterns

Y1: {s-M-d/d-M-d}-M-s



Y2: {s-M-d/d-M-d}-M-d



Data Structures **Algorithmic Modules**

s symbolic


d data

M inductive (ML)

K deductive (KR)

44 patterns were
discovered in systems
described by 476 papers


The Rise of Neuro-symbolic AI: AI System Patterns and their frequency



63% of systems use simple patterns such as:

- A1 (s-M-s) - link prediction on KG; rule learning
- (F) Fusion patterns that include both data and symbolic inputs

The Rise of Neuro-symbolic AI: Impact on Knowledge Engineering Systems




KE: creation and refinement of Knowledge Graphs, Ontologies and Taxonomies

KE tasks frequent among NeSy systems (25%)

A diversity of patterns used for KE tasks (18 from 44 total)


Frequent and specific patterns emerge

Which pattern to choose for which task?




Guides (junior) knowledge engineers in choosing “best practice” system architectures for a task at hand.

Knowledge Graph Completion



Knowledge Graph Completion: Through rule learning (T8)



AMIE+ (**ML**) system derives Horn rules (**sym2**) from large knowledge graphs (**sym1**).

When applied, they can lead to predictions (**sym4**) with precision up to 70%.

$y:isCitizenOf(x, y) \Rightarrow y:livesIn(x, y)$

$y:wasBornIn(x, y) \wedge y:isLocatedIn(y, z) \Rightarrow y:citizenOf(x, z)$

$y:hasWonPrize(x, G. W. Leibniz) \Rightarrow y:livesIn(x, Germany)$

$y:hasWonPrize(x, Grammy) \Rightarrow y:musicalRole(x, Guitar)$

$d:countySeat(x, y) \Rightarrow d:largestCity(x, y)$

$d:jurisdiction(z, y) \wedge d:successor(x, z) \Rightarrow d:jurisdiction(x, y)$


$w:ownedBy(x, y) \Rightarrow w:subsidiary(y, x)$

$w:relative(y, z) \wedge w:sister(z, x) \Rightarrow w:relative(x, y)$

y YAGO, w Wikidata, d DBpedia

L. Galárraga et al., Fast rule mining in ontological knowledge bases with AMIE+, VLDB Journal 24(6):707–730, 2015.


Knowledge Graph Completion: Through embeddings (I1)



Knowledge Graph

Embedded Representation

Machine Learning Task



- Graph embedding method (**ML1**) embeds a KG (**sym1**) into a vector space (d)
- Data (d) is further processed by a second ML component (**ML2**) to create a symbolic structure (e.g., links for link prediction, **sym2**).

Image source: By EdoardoRamalli - Own work, CC BY-SA 4.0,
<https://commons.wikimedia.org/w/index.php?curid=106644260>

Neuro-Symbolic AI – KG Completion in Pharmacology with I1





Journal of Biomedical Informatics

Volume 115, March 2021, 103696



Original Research

Drug repurposing for COVID-19 via
knowledge graph completion



R. Zhang et al., Drug repurposing for COVID-19 via knowledge graph completion, Journal of Biomedical Informatics, 2021.


Part 3: Challenges and emerging Paradigms



Knowledge Engineering Challenge 1

**Understand
interfaces between
symbolic and
sub-symbolic
components**

Knowledge
Engineering/
Semantic Web




Knowledge Graph Embeddings

Knowledge Graph

Embedded Representation

Machine Learning Task



bridge symbolic and subsymbolic AI by making (*symbolic*) KGs amenable to processing *with sub-symbolic* ML methods.


Important family of **Neuro-symbolic AI** systems (e.g., 37% of NeSys [1])

Key in **knowledge intensive** domains (medicine, finance, engineering) e.g., drug discovery

Image source: By EdoardoRamalli - Own work, CC BY-SA 4.0, <https://commons.wikimedia.org/w/index.php?curid=106644260>

[1] A. Breit et al. "Combining Machine Learning and Semantic Web: A Systematic Mapping Study". In: ACM Computing Surveys (2023).

Critical dissonance between Knowledge Engineering and KGEs!



(real-world) KGs are very diverse in terms of their *characteristics* (structure, connectivity, semantic richness)

YET, KGE methods are typically tested on a *few popular benchmarks* that do not reflect this diversity of characteristics [1]

The evaluation of KGE's primarily focused on non-KG-specific parameters (e.g., data splits, hyperparameter settings) [2]

Recent studies show that KG characteristics do play a key role in KGE performance

[1] R. Biswas et al. Knowledge Graph Embeddings: Open Challenges and Opportunities. TGDK, 2023.

[2] S. Bonner et al. Understanding the performance of knowledge graph embeddings in drug discovery. In: AI in the Life Sciences, 2022.



Charact.	Task	KG/Onto.	KGE	Conclusion	Ref.,Yr.
Node degree	Link prediction /Drug discovery	Hetionet	TransE/H, RotatE, ComplEX, DistMult	High-degree entities receive higher scores, revealing a structural bias that undermines prediction accuracy for low-degree nodes.	[9], 2022
Node degree, relationship frequency	Link prediction	FB15k(-237), WN18(RR)	TransE, DistMult	Low-degree nodes and low-frequency relations contain less information to be learned by embedding models, resulting in lower performance for link prediction.	[29], 2020
Node & attribute degree, relation exclusivity & co-occurrence, heterogeneity, negative sampling	Entity alignment	DBpedia, Wikidata, YAGO3 and their variations, movies and restaurants KGs	MTransE/(E+Rot), RDGCN, RREA, MultiKE, AttRE, KDCoE, BERT-INT	High density entities yield better performance; higher heterogeneity in descriptions and literals increases alignment difficulty; negative sampling helps the performance, unless the KGs are sparse.	[14], 2023
Tail in-degree, tail-relation degree.	Knowledge graph completion.	FB15k-237, NELL-995, CoDEx-M	ConvE, TuckER	Embeddings perform worse on triples with low-degree tail entities.	[34], 2023
Node centrality, relative position of entities	Link prediction	WN18RR, FB15k-237, ogbl-biokg	GFA-NN	Adding node centrality and position information into embeddings (GFA-NN model) improved link prediction, especially on graphs with high centrality variance.	[32], 2021
Semantic inclusion, node centrality	Link prediction	FB15k(237,238, 239)	TransE/R, RotatE, ComplEx, DistMult, RESCAL	Model-specific impact. TransE and RESCAL improve significantly; RotatE, ComplEx, and DistMult show marginal gains or reduced performance when adding semantic type and hierarchy relations.	[13], 2024
Lexical information, graph structure, logical constructors	Membership prediction, subsumption prediction.	HeLis, FoodOn, GO	OWL2Vec*	The use of textual information (e.g., entity labels) shows positive performance impact; Graph structure (inverse triples, axiom complexity) varies based on the ontology; logical axioms add useful semantics, esp. combined with the graph structure.	[12], 2021
Inverse relations	Link prediction	FB15K(-237), WN18(RR), YAGO3-10, Kinships.	RotatE ,MuRe, ConvE, ComplEX, TuckER, DistMult, SE QuatE, TransE	Explicitly modeling inverse relations improves performance of KGE models, especially with the local closed world assumption.	[2], 2022
Inverse, duplicate, cartesian product relations	Link prediction	FB15k, WN18, YAGO3-10	TransE/H/R/D, DistMult, ComplEx, ConvE, RotatE, TuckER	Inverse, duplicate and cartesian relations lead to inflated model performance by making link prediction easier and biasing embeddings toward learning inverse relation patterns	[1], 2020
SameAs, inverse relations	Recommender system	DBpedia, Wikidata, YAGO	RDF2Vec, TransE, DistMult	RDF2Vec benefited from sameAs relations; TransE&DistMult showed reduced performance. Inverse relations had no effect on TransE&DistMult, but negatively impacted RDF2Vec.	[17], 2024

E.g., *node degree bias* leads to better performance for well-connected entities [1]

BUT:
there is a lack of a systematic and comprehensive understanding of how KG characteristics affect the performance of KGE-enabled downstream tasks

[1] S. Bonner et al. *Implications of topological imbalance for representation learning on biomedical knowledge graphs*. In: *Briefings in Bioinf.* 23(5), 2022.

From Structures to Vectors: Decoding How **Knowledge Graph Characteristics** Shape **Embedding Strategies**



Which KG characteristics impact embedding performance, and how?

Do we need to rethink KG creation and evaluation?

What is a good KG for neuro-symbolic settings?


How to design KGE methods that consider KG characteristics?

How to build more diverse benchmarks that reflect real-life KGs?

Knowledge Engineering Challenge 2


**Re-evaluate and
re-design KE
methodologies**

Knowledge
Engineering/
Semantic Web



What are the KE Methodologies of the future?


Knowledge Engineering'2023



Knowledge Engineering'2033




What are the KE Methodologies of the future?




Fathallah, N., et al., NeOn-GPT: A Large Language Model-Powered Pipeline for Ontology Learning. ESWC'25.

Emerging NeSy Paradigm: Visual NeSy system patterns...



**...help identify emerging best
practices for building NeSy
systems in an area**


Emerging NeSy Paradigm: Visual NeSy system patterns...



**...support the design and
engineering of NeSy systems
through pattern based,
visual system engineering**


NeSy Perspective: Pattern-based engineering of AI Systems

Pattern-based engineering of AI Systems through visual and formal notation [1]




[1] Ekaputra, Fajar J. Pattern-based engineering of Neurosymbolic AI Systems. *J. of Web Semantics*, 2024.

NeSy Perspective: Pattern-based engineering of AI Systems




Bendler, J.E. et al. *Tool4Boxology: A Semantic Toolbox for Constructing and Analysing Neuro-Symbolic Architectures*. ESWC 2026.

Emerging NeSy Paradigm: Visual NeSy system patterns...



**... act as communication
artifacts among stakeholders
(i.e., AI ethics experts and
computer scientists)**

NeSy Perspective: Visual Patterns for Risk Elicitation




Patterns as communication support among diverse stakeholders:

- Knowledge engineers
- ML experts
- AI Ethics experts
- Project Managers

NeSy Perspective: Visual Patterns for Risk Elicitation: Example

GraphRAG-based Chatbot:
Retrieval of **event summaries based on user queries**



Pattern Notation based on: van Bekkum et al., Modular design patterns for hybrid learning and reasoning systems: a taxonomy, patterns and use cases. Applied Intelligence 51(9), 2021.

BEAM Visual Notation (Risk Perspective, adapted from [6])


[6] Golpayegani, D. et al., AIRO: an ontology for representing AI risks based on the proposed EU AI act and ISO risk management standards. In Towards a Knowledge-Aware AI, 2022.

System: Simple RAG-based Chatbot

The workflow describes a simplified event information chatbot based on the GraphRAG approach.

Context


The system will be used in a demonstration environment (not real-world use case) to showcase RAG systems for potential users. The Event KG is pre-populated with information about events for a chosen scenario (e.g., incident events in an organization)



Talk summary


Motivation

Knowledge is essential
in complex IS and
emerging technologies.




Neuro-symbolic KE


NeSy AI reshapes KE.



Challenges and Emerging Paradigms

We need to rethink KE
Patterns support NeSy
system engineering







NeSy AI ▼

Conferences ▼

 Login

NESY 2028

22nd INTERNATIONAL CONFERENCE ON NEUROSymbOLIC LEARNING AND REASONING

Vienna University of Economics and Business

September 13–15, 2028