

Trust Triggers'26:

Beyond Data: Why the Future of AI is Neurosymbolic



Marta Sabou



Beyond Data: Why the Future of AI is Neurosymbolic

PART 01



Motivation

Trust is not optional, and pure data-driven AI can't deliver it alone.

PART 02



The Emerging Approach

Neurosymbolic AI: an emerging field for enabling trustworthy AI systems.

PART 03



What's Missing?


How to engineer trustworthy neurosymbolic AI systems?

Purely Data-driven AI



GPT-4 generated

+ **AI** →



efficiency
flexibility
problem-solving



user
understanding
& trust



**What does an AI system need to know—
and be able to show—
to be trusted with real decisions?**

What trust requires: (Some) Dimensions of Trust



Explainability

Can it justify a decision in terms a person understands?



Verifiability


Can a claim be checked against known facts or rules?



Alignment

Does it respect domain knowledge, constraints, regulations?

Explaining Events in 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


Explanations of Events: Solution Concept based on KGs and Rules

ECar is charging slower (**E6**), because transformer is overloaded (**E4**), because houses consume more (**E5**, **E2**, **E3**) and produce less PV energy because of cloudy weather (**E1**)




Explainability Grounded in Domain Semantics

Explanations combine data-driven learning (events learned from sensor data) with explicit knowledge of the system — components, relations, causality rules



Compare to data-driven AI explanations: which pixels or features mattered.

A saliency heatmap



"These pixels mattered."



Explainability



Verifiability



Alignment

Eliciting and Exploiting Procedural Knowledge in Industry 5.0

Photo by Austrian National Library on
Unsplash



Procedural knowledge: knowing how, not just what

Industrial Procedural Knowledge is often **poorly digitalised and documented**, but the **correct execution of procedures** is crucial for the optimization of industry operations.

Hard to search, hard to transfer, hard to verify

Slide adapted from Irene Celino

Industrial Procedural Knowledge: BEKO Safety Procedure

LOTO: Lock-out Tag-out

- **list of tasks** to be **executed** in equipment management activities.
- ensures that **dangerous equipment** is **shut off** and not able to be started up **prior to the completion of maintenance** or repair work
- **stored as digital/paper structured documents** that should be **created** during first machine installation and **continuously updated** in case of machine modification
- multimodal: pictures, videos text.

ISTRUZIONE OPERATIVA DI SICUREZZA PER INTERVENTO E/O SBLOCCO
SULLA LINEA DI PREMONTAGGIO L3 - L4

Sito di Cassinetta Fabbrica Frigoriferi



Whirlpool
CORPORATION

ATTENZIONE

**PRIMA DI INIZIARE QUALSIASI TIPO DI ATTIVITA' DI MANUTENZIONE O DI INTERVENTO DI SBLOCCO/RIMOZIONE
SULLA LINEA**


Leggere attentamente questa istruzione e informare tutte le persone che possono essere interessate

DIFFERENTI TIPI DI ENERGIA CHE DEVONO ESSERE DISATTIVATI PRIMA DI QUALUNQUE INTERVENTO SULLA MACCHINA

ID	Tipo di energia	Processo di disconnessione	Verifica	Tipo di blocco	DPI da usare
	ELETTRICA	1 Agire sul quadro elettrico principale ruotando il sezionatore da I a O  Applicare il lucchetto ROSSO ed il TAG ROSSO	Riavviare la macchina per verificare L'ASSENZA di energia elettrica	lucchetto ROSSO e cartello ROSSO di identificazione (TAG) forniti in dotazione 	Guanti antitaglio Scarpe di sicurezza  
		Applicare il sistema di bloccaggio rosso per valvole e il TAG rosso 	Verificare che si sia scricata tutta l'aria 	Lucchetto e cartello (TAG) ROSSI sul ...	Guanti antitaglio Scarpe di sicurezza 

PERKS ambition and objectives: *Holistic Governance of Procedural Knowledge (PK)*


PERKS creates digital tools for elicitation and collection, management and access, reuse and exploitation of Procedural Knowledge in industrial settings



Slide adapted from Irene Celino

Industrial Procedural Knowledge: PERKS Architecture

Symbolic and subsymbolic AI are leveraged to build **Procedural Knowledge Graphs** to be incrementally populated and verified with **human-in-the-loop** approaches



Slide adapted from Irene Celino

Slide 12

M. SABOU - Beyond Data: Why the Future of AI is Neurosymbolic, TrustTriggers, Ghent'26

What trust requires: (Some) Dimensions of Trust



Explainability

Guidance through
procedure explained
based on structured
data.



Verifiability

Procedures can be
checked, not just
trusted blindly.



Alignment

Human-in-the-loop
validation: alignment with
how the plant actually
runs; PKG built with
workers, not just from
data

A Pattern is Emerging: Two projects, same approach



learning component + explicit knowledge component

This pairing of AI components is what turns the trust dimensions from wishes into an engineerable architecture.

Beyond Data: Why the Future of AI is Neurosymbolic

PART 01



Motivation

Trust is not optional, and pure data-driven AI can't deliver it alone.

PART 02



The Emerging Approach

Neurosymbolic AI: an emerging field for enabling trustworthy AI systems.

PART 03



What's Missing?

How to engineer trustworthy neurosymbolic AI systems?


Two AI Paradigms, Complementary Strengths

	Sub-symbolic AI (statistical learning from data)	Symbolic AI (explicit knowledge representation and reasoning)
Explainability	Weak: <i>Patterns, not justifications</i>	Strong: <i>Reasoning steps are inspectable</i>
Verifiability	Weak: <i>Confidence scores, not guarantees</i>	Strong: <i>Claims checked against rules or facts</i>
Alignment	Weak: <i>Context isn't natively represented</i>	Strong: <i>Domain knowledge is the substrate</i>
At scale	Strong: <i>Learns from noisy, unstructured data</i>	Weak: <i>Brittle, hard to scale, knowledge-acquisition bottleneck</i>

Trust needs both.


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

The Rise of Neuro-symbolic AI: AI System Patterns


Data Structures	Algorithmic Modules
s symbolic	M inductive (ML)
d data	K deductive (KR)
d data	{ begin parallel
s symbolic	/ separation parallel
M ML	} fusion parallel
K KR	- processing flow



d-M-s



{s-M-d/d}-M-s



highly complex pattern

AI system patterns [1] as abstractions to:

- reduce the complexity in understanding the internal workflow of AI system
- make AI system architecture transparent to stakeholders with heterogeneous AI knowledge

[1] F. van Harmelen, A. ten Teije, *A boxology of design patterns for hybrid learning and reasoning systems*, *J. of Web Engineering* 18 (2019) 97–124.

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](#)




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



44 patterns were
discovered in systems
described by 476 papers

The Rise of Neuro-symbolic AI: Not an anecdote: a field



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

*Sabou et al., Knowledge engineering in the
age of neurosymbolic systems, NAI 1, 2025.*




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.

The Neurosymbolic AI approach for trust is here.

476

Systems built on
neurosymbolic patterns

3

Provide provenance-
capturing mechanism

Beyond Data: Why the Future of AI is Neurosymbolic

PART 01



Motivation

Trust is not optional, and pure data-driven AI can't deliver it alone.

PART 02



The Emerging Approach

Neurosymbolic AI: an emerging field for enabling trustworthy AI systems.

PART 03



What's Missing?

How to engineer trustworthy neurosymbolic AI systems?

Auditing: Auditing as infrastructure, not afterthought


Many current AI system audits:

- have a narrow focus (primarily on system performance)
- are performed occasionally and manually
- are primarily “detective”



Continuous AI system auditing:


- has a broad focus on the entire AI system and its lifecycle
- uses continuous, automated monitoring
- is shifting towards “preventive”



Knowledge graphs as the audit trail: which data, which model, which rule fired?

Waltersdorfer et al., Leveraging Knowledge Graphs for AI System Auditing and Transparency, SWJ, 2025.

Auditing: The PERKS Approach




Slide adapted from Irene Celino

Slide 31

M. SABOU - Beyond Data: Why the Future of AI is Neurosymbolic, TrustTriggers, Ghent'26

The Human-in-the-loop: Who verifies, what and how?

- Neurosymbolic systems routinely rely on humans to check semantic resources for domain correctness, ontological consistency, or factual accuracy
- Complex area with still weak methodological foundations.



Sabou, Tsaneva, Fernandez, Poveda-Villalón & Suárez-Figueroa, *Human-centric Evaluation of Semantic Resources: A Systematic Mapping Study*, ACM Computing Surveys 58(10) (2026). DOI: 10.1145/3800939


The Human-in-the-loop: Human evaluation weakly documented



HUMAN-CENTRIC
ONTOLOGY
EVALUATION

WU
WIRTSCHAFTS
UNIVERSITÄT
WIEN VIENNA
UNIVERSITY OF
ECONOMICS
AND BUSINESS


The majority of evaluations do **not** report important aspects:



Sabou, Tsaneva, Fernandez, Poveda-Villalón & Suárez-Figueroa, *Human-centric Evaluation of Semantic Resources: A Systematic Mapping Study*, *ACM Computing Surveys* 58(10) (2026). DOI: 10.1145/3800939

The Human-in-the-loop: What are Suitable Human-machine workflows?


I. Human Judgement




II. AI Assistance



III. Human Verification




IV. Fully Automated



G. Faggioli, L. Dietz, C.L.A. Clarke, G. Demartini, M. Hagen, C. Hauff, et al. Perspectives on Large Language Models for Relevance Judgment. Proc. of the ACM SIGIR Int. Conf. on Theory of Information Retrieval (2023)

The Human-in-the-loop: What are Suitable Human-machine workflows?




Hybrid workflows of type II (AI Assistance) and III (Human verification) provide best results

Tsaneva, S., Dessì, D., Osborne, F., & Sabou, M. (2025). Knowledge graph validation by integrating LLMs and human-in-the-loop. *Information Processing & Management*, 62(5).

Systematic Risk Understanding: For Compliance with the EU AI Act

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.

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)




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.

Beyond Data: Why the Future of AI is Neurosymbolic


Beyond data:

*pure data-driven AI
alone can't deliver
trust.*




The future of AI is Neurosymbolic:

*NeSy enables trustworthy AI
systems by combining symbolic
and sub-symbolic AI methods.
It is not an anecdote, it is an
emerging field*



We are not there yet:

*Need to make NeSy AI
trustworthy themselves
(auditing, human-machine
workflows, alignment with
the EU AI Act)*





NeSy AI ▼

Conferences ▼

 Login

NESY 2028

22nd INTERNATIONAL CONFERENCE ON NEUROSYMBOLIC LEARNING AND REASONING

Vienna University of Economics and Business


September 13–15, 2028

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



What we do:

SemSys research landscape



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

Auditing: The PERKS Approach



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.