



Towards Neuro-Symbolic AI Applications at Bosch: Opportunities and Challenges

Daria Stepanova
Bosch Center for AI

My Background

- 2011 – 2015: PhD in computational logic at TU Vienna
 - Handling inconsistencies in combination of answer set programs and description logic ontologies
- 2015 – 2018: Postdoc and senior researcher at MPI for Informatics
 - Rule learning from knowledge graphs, shift to NeSy AI
- 2018 – present: Research scientist and then senior expert on hybrid AI at Bosch Center for AI
 - NeSy AI and its applications



About Bosch

Our Company in Figures

In 2024



417,900

Bosch associates
worldwide (approx.)

86,800

Associates work
in **R&D**, including

48,000
software developers



468

Subsidiaries and
regional companies in
60 countries

136

R&D locations
worldwide

Bosch Research Leveraging our International Setup



**Connect to
the best
academia**



**Connect to
RB business**



**Connect to
the local
economy**




**Live up to our
entrepreneurial
mission**



**Benefit from
local culture**



**Collaborate
with legal
frameworks**



Renningen Research Campus

Overview



**Bosch associates on site
(including PhD)**

~2,100

5 of which from Corporate Research alone

~1,400

Who we are

Our Business Sectors



Mobility Solutions



Industrial Technology



**Energy and Building
Technology**



Consumer Goods

AI at Bosch

Bosch vs Other AI Companies



```
...
def __init__(self, threshold):
    self.threshold = threshold

def __getitem__(self, item):
    if item not in self.items:
        raise ValueError(f'Item {item} not in the dictionary')
    return self.items[item]

def __setitem__(self, key, value):
    self.items[key] = value

def __delitem__(self, key):
    del self.items[key]

def __contains__(self, key):
    return key in self.items

def __iter__(self):
    return iter(self.items)

def __len__(self):
    return len(self.items)

def __str__(self):
    return str(self.items)

def __repr__(self):
    return repr(self.items)

def __eq__(self, other):
    return self.items == other.items

def __neq__(self, other):
    return self.items != other.items

def __lt__(self, other):
    return self.items < other.items

def __gt__(self, other):
    return self.items > other.items

def __le__(self, other):
    return self.items <= other.items

def __ge__(self, other):
    return self.items >= other.items

def __add__(self, other):
    return self.items + other.items

def __sub__(self, other):
    return self.items - other.items

def __mul__(self, other):
    return self.items * other.items

def __div__(self, other):
    return self.items / other.items

def __mod__(self, other):
    return self.items % other.items

def __pow__(self, other):
    return self.items ** other.items

def __call__(self, *args, **kwargs):
    return self.items(*args, **kwargs)

def __getattr__(self, name):
    return getattr(self.items, name)

def __setattr__(self, name, value):
    setattr(self.items, name, value)

def __delattr__(self, name):
    delattr(self.items, name)

def __copy__(self):
    return self.__class__(self.items)

def __deepcopy__(self, memo):
    return self.__class__(self.items)

def __reduce__(self):
    return (self.__class__, (self.items,))

def __reduce_ex__(self, protocol):
    return (self.__class__, (self.items,))

def __sizeof__(self):
    return self.items.__sizeof__()

def __dir__(self):
    return dir(self.items)


def __hash__(self):
    return hash(self.items)

def __bool__(self):
    return bool(self.items)

def __bool__(self):
    return bool(self.items)

def __bool__(self):
    return bool(self.items)

```



1886



1950



2025



AI at Bosch

Bosch vs Other AI Companies

product design
factories
quality control
manufacturing
hardware assembly lines
warehousing
production scheduling
manual inspection
engineering
supply chains
simulation
logistics
physics

machine learning
foundation models
optimization
algorithms
model deployment
software
computer vision
data
data pipelines
GenAI
neural networks
natural language processing
pattern recognition
reinforcement learning

 **BOSCH**



AI



1886

1950

2025

AI at Bosch



digital twins

domain knowledge

sensor data

industrial AI

smart manufacturing

embedded AI anomaly detection

AI at  **BOSCH**

 **BOSCH**

1886

learning
optimization
hms
computer vision
data pipelines
data
neural networks
natural language processing
cognition
reinforcement learning



2025

Who we are

Bosch Center for AI

Founded in
2017

Major Research Activities

- Generative AI
- Large Language Models
- Prediction and Planning
- Verification and Validation of AI Models
- **Neuro-symbolic AI**
- Multi-modal Foundation Models
- Time-series Foundation Models

Application Areas


- Automated Driving, Driving Assistants
- Industrial Robotics
- Knowledge Management
- Smart Consumer Goods
- AI for Engineering & Manufacturing
- Embedded AI
- Manufacturing

Impact

- #1 AI research lab in the European industry
- Since 2017*: 350+ top-tier publications in AI
- 1700+ patent filings*

Research Fields and Core Competences


Deep Neural Networks for Automated Driving perception



Manufacturing optimization, e.g., optical inspection



System engineering



Probabilistic Modeling



Deep Learning



NLP & Neuro-symbolic AI



RL & Control



Collaborations



Carnegie Mellon University



* Status as of 2024

Artificial Intelligence @ Bosch

A lot of Promises

2021



“

*By **2025** the aim is for all Bosch products to either contain AI or have been developed or manufactured with its help.*

Bosch Vision for AI

*Volkmar Denner
Bosch CEO 2012-2021*

Artificial Intelligence @ Bosch

A lot of Expectations

2023




“


*This year we will achieve our goal **ahead of schedule** that all Bosch products and solutions contain AI or were developed or produced with AI.*

Stefan Hartung
Bosch CEO, Handelsblatt 14. April 2023

Symbolic AI




A lot of valuable
domain knowledge




Challenges: Modeling efforts and scalability

Neural AI



Vast amounts
of data



Challenges: Explainability and trustworthiness

Neuro-symbolic AI at Bosch




Pros:

Explainability
Trustworthiness

Cons:

Modeling efforts
Lack of scalability



Cons:

Lack of explainability
Lack of trustworthiness

Pros:

No modeling efforts
Scalability

Research questions

Q1: How to automatically curate and complete KGs, learn rules?

Q2: How to use KGs, rules to improve/explain ML models?

Q3: How to scale up and democratize reasoning systems for industry?

Q4: Industrial innovations with neuro-symbolic AI?

Neuro-symbolic AI at Bosch




Pros:

Explainability
Trustworthiness

Cons:

Modeling efforts
Lack of scalability



Cons:

Lack of explainability
Lack of trustworthiness

Pros:

No modeling efforts
Scalability

Research questions

Q1: How to automatically curate and complete KGs, learn rules?

Q2: How to use KGs, rules to improve/explain ML models?

Q3: How to scale up and democratize reasoning systems for industry?

Q4: Industrial innovations with neuro-symbolic AI?

Neuro-symbolic AI at Bosch




Pros:

Explainability
Trustworthiness

Cons:

Modeling efforts
Lack of scalability



Cons:

Lack of explainability
Lack of trustworthiness

Pros:

No modeling efforts
Scalability

Research questions

Q1: How to automatically curate and complete KGs, learn rules?

Q2: How to use KGs, rules to improve/explain ML models?

Q3: How to scale up and democratize reasoning systems for industry?

Q4: Industrial innovations with neuro-symbolic AI?

Neuro-symbolic AI at Bosch




Pros:

Explainability
Trustworthiness

Cons:

Modeling efforts
Lack of scalability



Cons:

Lack of explainability
Lack of trustworthiness

Pros:

No modeling efforts
Scalability

Research questions

Q1: How to automatically curate and complete KGs, learn rules?

Q2: How to use KGs, rules to improve/explain ML models?

Q3: How to scale up and democratize reasoning systems for industry?

Q4: Industrial innovations with neuro-symbolic AI?

Neuro-symbolic AI at Bosch




Pros:

Explainability
Trustworthiness

Cons:

Modeling efforts
Lack of scalability



Cons:

Lack of explainability
Lack of trustworthiness

Pros:

No modeling efforts
Scalability

Research questions

Q1: How to automatically curate and complete KGs, learn rules?

Q2: How to use KGs, rules to improve/explain ML models?

Q3: How to scale up and democratize reasoning systems for industry?

Q4: Industrial innovations with neuro-symbolic AI?


Neuro-symbolic AI at Bosch

Pros:

Explainability
Trustworthiness

Cons:

Modeling efforts
Lack of scalability



Cons:

Lack of explainability
Lack of trustworthiness

Pros:

No modeling efforts
Scalability

Research questions

Q1: How to automatically curate and complete KGs, learn rules?

Q2: How to use KGs, rules to improve/explain ML models?

Q3: How to scale up and democratize reasoning systems for industry?

Q4: Industrial innovations with neuro-symbolic AI?

GenAI?

Talk Overview

- NeSy AI at Bosch
- Knowledge graph curation and completion
 - Market analysis
- Scaling and democratizing reasoning systems
 - Machine scheduling
 - E-drive configuration
- Explaining ML models
 - Edge case analysis in autonomous driving
- Summary
 - Opportunities
 - Challenges



Talk Overview

- NeSy AI at Bosch
- **Knowledge graph curation and completion**
 - **Market analysis**
- Scaling and democratizing reasoning systems
 - Machine scheduling
 - E-drive configuration
- Explaining ML models
 - Edge case analysis in autonomous driving
- Summary
 - Opportunities
 - Challenges

Market Analysis



- Predict potential competitors
- Suggest potential suppliers
- Optimize number of suppliers and their carbon footprint
- Analyze business sectors in which companies along supplier chain are involved




Numerical Values are Important

1. Extend KG with **numerical information**

- Carbon footprint
- Average revenue
- Number of employees
- ...


Information extraction
to mine numerical values
for entities from the web



KG Incompleteness

2. Learn rules from KGs with numericals, apply them to predict facts, e.g., potential suppliers

Extension of NeuralLP
[Yang et al, NeurIPS 2017]
for rule extraction to
handle numerical values



$$\text{hasSupplier}(X, Z) \leftarrow \text{hasSupplier}(X, Y), \text{hasCompetitor}(Y, Z), \\ \text{hasRevenue}(Z, N1), \text{hasRevenue}(Y, N2), N1 > N2$$

KG Incompleteness

3. Avoid nonsensical predictions, e.g., company has city as a supplier

Our reasoning-based method

Standard embedding methods




N. Jain, T.K. Tran, M. Gad-Elrab, D. Stepanova: Improving Knowledge Graph Embeddings with Ontological Reasoning. ISWC 2021

M. Andresel, T.K. Tran, M. C. Domokos, P. Minervini, D. Stepanova: Combining Inductive and deductive Reasoning for Query Answering over Incomplete KGs. CIKM 2023

Explainable Entity Similarity

4. Company similarity: Given a supplier, suggest similar ones, and explain why they are similar



Y. Ismaeil, D. Stepanova, K. T. Tran, H. Bloeckel. FEABl: A Feature Selection-Based Framework for Interpreting KG Embeddings. ISWC 2022.


D. Daza, C. Xuan Chu, T.K. Tran, D. Stepanova, M. Cochez, P. Groth: Explaining Graph Neural Networks for Node Similarity on Graphs. TLMR 2026.

See tutorial <https://essai2026.eu/program.php#id-T8> on KG embeddings

Supplier Optimization

Simplified Real Example

11 potential suppliers




Supplier Optimization

Simplified Real Example


11 potential suppliers

4 current suppliers




Supplier Optimization

Simplified Real Example



Supplier Optimization

Simplified Real Example



Talk Overview

- NeSy AI at Bosch
- Knowledge graph curation and completion
 - Market analysis
- **Democratizing and scaling reasoning systems**
 - **Machine scheduling**
 - **E-drive configuration**
- Explaining ML models
 - Edge case analysis in autonomous driving
- Summary
 - Opportunities
 - Challenges

LLMs for Reasoning Problems?

I need to seat 7 people around a table at my wedding reception, but my parents and in-laws should not sit next to each other. Also, my wife insists we look better in pictures when she's on my left, but I need to sit next to my best man. How do I seat us on a round table? But then, what happens if we invited our pastor to sit with us?

Send

Traditional LLM Model

Tokens: 439



Reasoning Model


Tokens: 8,559

Reasoning On



20x More Tokens
150x More Compute

0:34 / 0:42 • Human Touch



Reasoning LLMs for Combinatorial Problems


Theoretical and Practical Limitations

QUANTIFYING GENERALIZATION COMPLEXITY FOR LARGE LANGUAGE MODELS

Zhenting Qi^{1*}, Hongyin Luo², Xuliang Huang³, Zhuokai Zhao^{4,5}, Yibo Jiang⁵

Xiangjun Fan⁴, Himabindu Lakkaraju¹, James Glass²
¹Harvard University, ²Massachusetts Institute of Technology

³University of Illinois at Urbana-Champaign, ⁴Meta, ⁵University of Chicago
zhentingqi@u.illinois.edu, hyluo@mit.edu



ICLR 2025

On Limitations of the Transformer Architecture

Binghui Peng
Department of Computer Science
Columbia University
bp2601@columbia.edu

Srini Narayanan
Google DeepMind
srinin@google.com

Christos Papadimitriou
Department of Computer Science
Columbia University
christos@columbia.edu



COLM 2024


NPHardEval: Dynamic Benchmark on Reasoning Ability of Large Language Models via Complexity Classes

Lizhou Fan^{1*} Wenyue Hua^{2*} Lingyao Li¹ Haoyang Ling¹ Yongfeng Zhang²

¹School of Information, University of Michigan, Ann Arbor, MI 48103

²Department of Computer Science, Rutgers University, New Brunswick, NJ 08854
{lizhouf, lingyao, hvfrankl}@umich.edu, {wenvue.hua, yongfeng.zhang}@rutgers.edu

*Lizhou Fan and Wenyue Hua contribute equally.



ACL 2024


Have Large Language Models Learned to Reason? A Characterization via 3-SAT Phase Transition

Rishi Hazra¹, Gabriele Venturato², Pedro Zuidberg Dos Martires¹ & Luc De Raedt^{1,2}¹Centre for Applied Autonomous Sensor Systems (AASS),


Örebro University, 70182, Sweden

²Department of Computer Science, KU Leuven, 3001, Belgium

{rishi.hazra, pedro.zuidberg-dos-martires}@oru.se
{gabriele.venturato, luc.deraedt}@kuleuven.be



COLM 2025



ICLR 2026

Limits of Transformer Language Models on Learning to Compose Algorithms

Jonathan Thomm^{1,2*}
jthomm@ethz.ch

Giacomo Camposampiero^{1,2}
giacomo.camposampiero1@ibm.com

Aleksandar Terzic^{1,2}
aleksandar.terzic1@ibm.com

Michael Hersche¹
michael.hersche@ibm.

Bernhard Schölkopf^{2,3}
bs@tuebingen.mpg.de

Abbas Rahimi¹
abr@zurich.ibm.com

¹IBM Research – Zurich, ²ETH Zurich, ³MPI Tübingen

FOCS 2025

NeurIPS 2024

Answer Set Programming (ASP)

Overview



Thomas
Eiter



Nysret
Musliu




Johannes
Oetsch




Tobias
Geibinger



Nelson
Higuera



Tobias
Kaminski




BOSCH

$$\text{ASP} = \text{DB} + \text{KR} + \text{SAT/SMT}$$

See tutorial <https://essai2026.eu/program.php#id-T2> for more details on ASP


Answer Set Programming (ASP)

Example



Answer Set Programming (ASP)


Example



Answer Set Programming (ASP)

Application Areas at Bosch

In collaboration with




Challenge 1: Scalability for Industry

Parallel Machine Scheduling




- Challenging parallel machine scheduling problem at semiconductor production plant
 - sequence-dependent setup times
 - release dates
 - machine capabilities
- **Goal:** maximise throughput (number of jobs / time unit)
 - minimal makespan (total schedule length) achieves this, but refined lexicographic minimal makespan is better
- **Solution:** extension of ASP with difference logic, clingo-dl solver used for computing results

In collaboration with




Lexicographic makespan



T. Eiter, T. Geibinger, N. Musliu, J. Oetsch, P. Skocovský, D. Stepanova: Answer-Set Programming for Lexicographical Makespan Optimisation in Parallel Machine Scheduling. *Theory Pract. Log. Program.* 23(6): 1281-1306 (2023)

Challenge 1: Scalability for Industry


Adaptive Large Neighborhood Search for ASP Optimization



T. Eiter, T. Geibinger, N. Higuera Ruiz, N. Musliu, J. Oetsch, D. Pfliegler, D. Stepanova: Adaptive large-neighbourhood search for optimisation in answer-set programming. *Artif. Intell.* 337: 104230 (2024)

Challenge 2: Interaction with ASP Solvers for Non-Experts


E-Drive Conceptual Design and Configuration



Iterate on requirements

Challenge 2: Interaction with ASP Solvers for Non-Experts

Combining ASP and LLMs




[Eiter *et al*, IJCAI 2023, Eiter *et al*, JELIA 2023, Geibinger *et al*, TAASP 2024]

Challenge 2: Interaction with ASP Solvers for Non-Experts

Combining ASP and LLMs


- ASP + LLM for explaining solutions
- Questions from users on initial solution:
 - Why **is the torque density high**?
 - Why **is the medium air** and **not water**?
 - Why is the voltage **not low**?
 - Why is there **no** design found fulfilling requirements?
- Translate question into counterfactual which does not hold, employing LLMs
- Explain why counterfactual is inconsistent [Herud et al, ICLP 2022]

[Eiter et al, IJCAI 2023, Eiter et al, JELIA 2023, Geibinger et al, TAASP 2024]



Challenge 3: Automating Modeling Efforts

Combining ASP and LLMs



Open Challenge: Ensure encoding correctness without revealing the ASP code to the user

[Schrader *et al*, AAAI 2026]

Neuro-symbolic Reasoning

Trends and Perspectives for Bosch





Image from <https://arxiv.org/pdf/2505.24874>

Neuro-symbolic Reasoning

Trends and Perspectives for Bosch

Logic-based NeSy frameworks according to the 6 dimensions outlined in the paper.

Frameworks	Inference	Syntax	Semantics	Learning	Representations	Paradigms	Tasks
	(P)roof (M)odel	(P)ropositional (R)elational (C)lassical (F)uzzy (P)robability	(M)inimal (S)table (C)lassical (F)uzzy (P)robability	(P)arameters (S)tructure	(S)ymbolic (Sub)symbolic	Logic (L/I) Probability (P/p) Neural(N/n)	(D)istant (S)upervision (S)emi (S)upervised (KGC)ompletion (G)enerative (K)nowledge (I)nduction
α ILP [111]	P+M	FOL	S + P	P + S	S	Ln	KI
∂ ILP [39]	P	R	M + F	P + S	S	Ln	DS + KI
DeepProbLog [72]	P+M	FOL	M + P	P+S	S+Sub	LpN	DS + KI
DeepStochLog [132]	P	FOL	M + P	P	S	LpN	DS + SS
DiffLog [112]	P	R	M + F	P+S	S	Ln	KI
DL2 [40]	M	P	C + F	P	S+Sub	IN	DS + SS
DLM [77]	M	FOL	C + F + P	P	S	IPN	SS + KGC
LRNN [116]	P	R	M + F	P + S	S + Sub	LN	KGC + KI
LTN [5]	M	FOL	C + F	P	S + Sub	IN	DS + SS
NeuralLP [137]	P	R	M + F	P	S	Ln	KGC + KI
NeurASP [138]	P+M	FOL	S + P	P	S	LpN	DS
NLM [35]	P	R	M + F	P + S	S	Ln	KGC + KI
NLog [121]	P	R	M + P	P	S	LpN	DS
NLProlog [131]	P	R	M + P	P + S	S + Sub	LpN	KGC + KI
NMLN [78]	M	FOL	C + P	P + S	S + Sub	IPN	KGC + G
NTP [102]	P	R	M + F	P + S	S + Sub	Ln	KGC + KI
RNM [76]	M	FOL	C + P	P	S	IPN	SS
SBR [33]	M	FOL	C + F	P	S+Sub	IN	DS + SS
Scallop [58]	P	FOL	M + P	P	S	LpN	DS
SL [133]	M	P	C + P	P	S	LpN	SS
Slash [113]	P+M	FOL	S + P	P	S	LpN	DS + SS
TensorLog [18]	P	R	M + P	P	S	LpN	DS + KGC



[Marra *et al*, AIJ 2024]

Image from <https://arxiv.org/pdf/2505.24874>

Need for established industrial frameworks

See tutorial <https://essai2026.eu/program.php#id-T6> on foundations of NeSy

Neuro-symbolic Reasoning

Trends and Perspectives for Bosch

See tutorial <https://essai2026.eu/program.php#id-T5> on MCP-Solver

MCP-Solver: Integrating Language Models with
Constraint Programming Systems

Stefan Szeider

Algorithms and Complexity Group
TU Wien, Vienna, Austria
www.ac.tuwien.ac.at/people/szeider/

Logic-based NeSy frameworks according to the 6 dimensions outlined in the paper.

Frameworks	Inference	Syntax	Semantics	Learning	Representations	Paradigms	Tasks
	(P)roof (M)odel	(P)ropositional (R)elational (F)OL	(M)inimal (S)table (C)lassical (P)uzzy (P)robability	(P)arameters (S)tructure	(S)ymbolic (S)ubsymbolic	Logic (L/I) Probability (P/p) Neural(N/n)	(D)istant (S)upervision (S)emi (S)upervised (KGC)ompletion (G)enerative (K)nowledge (I)nduction
α ILP [111]	P+M	FOL	S + P	P + S	S	Ln	KI
∂ ILP [39]	P	R	M + F	P + S	S	Ln	DS + KI
DeepProbLog [72]	P+M	FOL	M + P	P+S	S+Sub	LpN	DS + KI
DeepStochLog [132]	P	FOL	M + P	P	S	LpN	DS + SS
DiffLog [112]	P	R	M + F	P+S	S	Ln	KI
DL2 [40]	M	P	C + F	P	S+Sub	IN	DS + SS
DLM [77]	M	FOL	C + F + P	P	S	IPN	SS + KGC
LRNN [116]	P	R	M + F	P + S	S + Sub	LN	KGC + KI
LTN [5]	M	FOL	C + F	P	S + Sub	IN	DS + SS
NeuralLP [137]	P	R	M + F	P	S	Ln	KGC + KI
NeuralASP [138]	P+M	FOL	S + P	P	S	LpN	DS
NLM [35]	P	R	M + F	P + S	S	Ln	KGC + KI
NLog [121]	P	R	M + P	P	S	LpN	DS
NLProlog [131]	P	R	M + P	P + S	S + Sub	LpN	KGC + KI
NMLN [78]	M	FOL	C + P	P + S	S + Sub	IPN	KGC + G
NTP [102]	P	R	M + F	P + S	S + Sub	Ln	KGC + KI
RNM [76]	M	FOL	C + P	P	S	IPN	SS
SBR [33]	M	FOL	C + F	P	S+Sub	IN	DS + SS
Scallop [58]	P	FOL	M + P	P	S	LpN	DS
SL [133]	M	P	C + P	P	S	LpN	SS
Slash [113]	P+M	FOL	S + P	P	S	LpN	DS + SS
TensorLog [18]	P	R	M + P	P	S	LpN	DS + KGC

[Marra et al, AIJ 2024]

Need for established industrial frameworks





Image from <https://arxiv.org/pdf/2505.24874>

See tutorial <https://essai2026.eu/program.php#id-T6> on foundations of NeSy

Talk Overview

- NeSy AI at Bosch
- Knowledge graph curation and completion
 - Market analysis
- Democratizing and scaling reasoning systems
 - Machine scheduling
 - E-drive configuration
- **Explaining ML models**
 - **Edge case analysis in autonomous driving**
- Summary
 - Opportunities
 - Challenges

Edge Cases in Autonomous Driving




How a tesla misinterprets an old carriage

<https://x.com/RealRusty/status/1559636816101244931>



Systematic Error Detection for Edge Case Generation

Motivation & Challenge



Images from Bosch colleagues

- Data collection is time-consuming, particularly for corner cases.
- Scarcity of data often limits fast development and release of AI models.

Common data



Corner cases



Real corner cases are challenging to collect at large-scale

<https://www.intc.com/news-events/press-releases/detail/1518/mobileyes-self-driving-secret-200pb-of-data>

Identifying Systematic Errors of Image Classifiers

What is a Systematic Error

- Subset of the data (“slice”) on which a given model underperforms (high error)
- Slices need to be interpretable by humans
 - “hopping kangaroos”
 - “white trailer + bright sky”
 - “people with dark skin tone”
- Systematic error identification is the task of finding such slices
 - Ideally also without required labelled data covering these rare slices
- Relevant for safety, reliability and fairness considerations in AI-based systems

Volvo admits its self-driving cars are confused by kangaroos

Swedish company's animal detection system can identify and avoid deer, elk and caribou, but is yet to work against the marsupials' movements



📷 Kangaroos are responsible for about 90% of collisions between vehicles and animals in Australia – although most are not serious. Photograph: Paul Kane/Getty Images

See tutorial <https://essai2026.eu/program.php#id-C7>
on need for AI explainability and trustworthiness

Identifying Systematic Errors of Image Classifiers

Overview

User specified input

Viewpoint: {front, rear, side..}


Color: {red, green, orange..}

Weather: {sunny, rainy, snowy..}

...

Class: minivan

viewpoint : rear
size : small
color : orange
weather : snowy
background: forest
ConvNeXt-B vs. VGG16



...

viewpoint : front
size : large
color : green
weather : lightning
background: house
ViT-B/16 vs. ResNet50



Combinatorial test cases for an image classifier generated

viewpoint : front
size :
color :
weather : sunny
background: people
ResNet50 vs. ViT-L/32



[Metzen *et al*, ICCV 2023]

Identifying Systematic Errors of Image Classifiers

Prompt Attack

[Metzen *et al*, ICCV 2023]

Test engineer specifies:

Source Class, Operational Design Domain, Prompt Template, and Risk Assessment

Operational Design Domain $\mathbf{Z}$

$\mathbf{Z}_0$ viewpoint: {front, rear, side,..}

$\mathbf{Z}_1$ color: {red, green, orange,...}

$\mathbf{Z}_2$ weather: {sunny, rainy, snowy, ...}

$\mathbf{Z}_3$ background: {lake, house, tree, ...}

Prompt Template

"{viewpoint} view of {color}
{source_class} in front of
{weather} {background}."

T_p

Source Class

source_class: minivan

$\tilde{\mathbf{y}}$

Risk Assessment

$\text{median}_{i=0}^{n_s-1} (1 - f(\tilde{\mathbf{y}}|\mathbf{x}_i))$

ranked
test cases


Identifying Systematic Errors of Image Classifiers

Prompt Attack

[Metzen *et al*, ICCV 2023]


PromptAttack handles:

Combinatorial Testing, Image Synthesis, Risk-based Ranking



Identifying Systematic Errors of Image Classifiers


Avoid Manual Labels



Y. Ismaeil, J-H. Metzen, K. T. Tran, H. Bloeckel, D. Stepanova. Beyond Manual Labels: Unsupervised Graph-Based Explanations for Error Analysis in Image Classifiers. ISWC 2025


Identifying Systematic Errors of Image Classifiers

Avoid Manual Labels



ILP for rule extraction in
[Collevati *et al*, NeSy 2024]

Vision-language models for detailed
scene graph construction



cgSpan [Schaul *et al*, IEEE 2021] for frequent
graph pattern mining over scene graphs

Output: *Images with ginger dogs and
black cars are frequently wrongly classified*

Talk Overview


- NeSy AI at Bosch
- Knowledge graph curation and completion
 - Market analysis
- Democratizing and scaling reasoning systems
 - Machine scheduling
 - E-drive configuration
- Explaining ML models
 - Edge case analysis in autonomous driving

▪ **Summary and outlook**

- **Opportunities**
- **Challenges**

Summary


Long-term agenda on automatic KG curation



Populating KGs with numerical facts WWW'22




Ontologies and embeddings for QA over KGs CIKM'23




Learning rules from KGs using embeddings CIKM'23, ICLR'20

Ontology-enhanced embeddings ISWC'21



Explainable clustering over KG embeddings ISWC'20



Explaining KG embeddings ISWC'23

Industrial applications



Market analysis







Autonomous driving




Manufacturing

Scaling and democratizing reasoning systems




Recent works on explainability of neural networks using KGs



Combining LLMs and KGs for detecting bias in image classifiers ISWC'25


Explaining node similarity in GNNs TLMR'26



Interpreting neurons using KGs ISWC'22

Large neighborhood search for improving scalability of answer set programming for parallel machine scheduling, e-drive configuration, supplier optimization [AIJ'24](#), [TPLP'23](#), [AAAI'22](#), [KR'22](#), [KR'21](#), ASP+LLMs

Other NeSy AI Works at Bosch





Bosch Semantic Stack Documentation

Bosch Semantic Stack master

Search




Bosch Semantic Stack > Concepts



The screenshot shows the TopBraid EDG Enterprise Data Governance interface. It features a 'Welcome to the STEADY production server!' message. Below this is a 'Recently Viewed Collections' section with a list of collections. The 'Asset Collections Summary' section shows a table of collections, with '652 - Ontologies' highlighted in red. The 'My Governance Roles' section shows a table of roles and their corresponding collections. The 'Workflow Summary' section shows a list of workflows. The 'My Tasks' section shows a list of tasks. The 'What's Happening?' section shows a list of events.

Other NeSy AI Works at Bosch



Bosch Semantic Stack Documentation

Bosch Semantic Stack > Concepts


Solutions



Bosch Semantic Stack

Digital Twin Registry



Aspect Model Catalog



Heterogenous Data Sources



- Ontology learning from text
- GraphRAG
- Knowledge-enriched document labeling [Thin *et al*, EMNLP 2025]
- Semantic mapping of terminologies [Schmid *et al*, ESWC 2025]
- Knowledge-enriched visual anomaly detection [Sadikaj *et al*, ICCV 2025]
- KG-enriched ML pipelines [Klironomous *et al*, ESWC 2023]
- Formal software verification + LLMs




Outlook: Promising NeSy AI Applications

Hardware Design


User Design Specification

1. Design a high voltage high power converter w/ MCU and sensors as controller, add another ROM as external storage
2. Chip Vendor: Ti, ADI, NXP, etc.

Customer requirements



Schematics



Layout




Printed Circuit Board product

Outlook: Promising NeSy AI Applications

Hardware Design

Neural generation of designs




Symbolic verification

PCBSchemaGen [Zou *et al*, Arxiv 2026], also PCBGPT [King *et al*, 2026]

Visual Question Answering


Schematic Images

Object detection and counting



Q: How many large things are either cyan metallic cylinders or yellow blocks?

[CLEVR, Johnson *et al*, CVPR 2017]



What is the total component count in this schematic?

Claude-Opus



7



[Lu *et al*, ECCV 2026]
Omnisch.com

Visual Question Answering Schematic Images

Topological reasoning




Are Pring Square and Rolv Grove connected?
[Bauer *et al*, ICLP 2024]

[Lu *et al*, ECCV 2026]
Omnisch.com


Visual Question Answering

Schematic Images

Spatial reasoning



A: Complete scene




B: Incomplete scene

There is a blue rubber object that is in front of the small metallic cube and right of the large object; what shape is it?

C: Question about hidden object

CLEVR-POC [Abraham *et al*, LREC/COLING 2024]

- Electronics background knowledge
- Graph-based reasoning
- Occluded symbols (e.g., in scans)



List the symbol names of the components right side of Q3.

QwenVL

[U2, C8, C2]




[Lu *et al*, ECCV 2026]

Omnisch.com

Neuro-Symbolic AI at Bosch

Opportunities

- Vast amounts of **domain expert knowledge** in various forms
 - Electronic engineers, production experts, material scientists
 - Knowledge is often in heads of people, but also text, csv files, etc.
- Availability of **large amounts of data**
 - Time series, production data, collected images
- Many **applications** where generative AI/ML alone is insufficient
 - Logistics, manufacturing optimization, system engineering



Neuro-Symbolic AI at Bosch

Challenges


- **Concrete value:** Hard to demonstrate it in industry, (not only for neuro-symbolic AI, but AI in general)
- **Hype:** Focus of the company and investments have been exclusively in GenAI for the past couple of years
 - Neuro-symbolic AI has been often questioned:
 - What is it exactly, why needed, concrete benefits, etc.
- **Lack of expertise in business units:**
 - Not many understand what rules, ontologies, reasoners are
 - A lot of efforts are still spent on “spreading the word” about symbolic AI
- **Modeling bottleneck:** for non-experts encoding knowledge and using symbolic reasoners is much more difficult than prompting a GenAI model

Neuro-Symbolic AI at Bosch

Outlook

- Prompt-symbolic focus mainly so far:
 - Ontology construction with GenAI, GraphRAG
 - For advanced QA only or beyond?
E.g., inconsistency handling of requirements?
 - Translation of natural language text to formal rules/constraints
 - E.g., MCP-Solver [Szeider, 2025]
 - Utilizing symbolic AI tools (e.g., ASP) while hiding users from encodings
(is it at all possible / to which degree?)
- Enriching tabular foundation models with contextual knowledge, e.g., [Boshko *et al*, Arxiv 2026]
- Applications in core engineering domains

How to leverage NeSy soup in industry?



Logic-based NeSy frameworks according to the 6 dimensions outlined in the paper.

Frameworks	Inference	Syntax	Semantics	Learning	Representations	Paradigms	Tasks
	(Proof Model)	(Propositional Relational FOL)	(Minimal Stable Classical Fuzzy Probability)	(Parameters Structure)	(Symbolic Subsymbolic)	Logic (L/I) Probability (P/p) Neural(N/n)	(Distance Supervision) (Semi Supervised) (KGCompletion) (Generative) (Knowledge Induction)
oLP [111]	P+M	FOL	S + P	P + S	S	Ln	KI
dLP [99]	P	R	M + F	P + S	S	Ln	DS + KI
DeepProbLog [72]	P+M	FOL	M + P	P+S	S+Sub	LpN	DS + KI
DeepStochLog [132]	P	FOL	M + P	P	S	LpN	DS + SS
DiffLog [112]	P	R	M + F	P+S	S	KI	DS + SS
DL2 [61]	P	C + F	P	S+Sub	IN	DS + SS	DS + SS
DLM [77]	M	FOL	C + F + P	P	S	IPN	SS + KGC
LRNN [116]	P	R	M + F	P + S	S + Sub	LN	KGC + KI
LTR [5]	M	FOL	C + F	P	S + Sub	IN	DS + SS
NeuralLP [137]	P	R	M + F	P	S	Ln	KGC + KI
NeuralSP [138]	P+M	FOL	S + P	P	S	LpN	DS
NLM [53]	P	R	M + F	P + S	S	Ln	KGC + KI
NLog [121]	P	R	M + P	P	S	LpN	DS
NLogProlog [131]	P	R	M + P	P + S	S + Sub	LpN	KGC + KI
NLMN [78]	M	FOL	C + P	P + S	S + Sub	IPN	KGC + G
NTP [102]	P	R	M + F	P + S	S + Sub	Ln	KGC + KI
RNN [76]	M	FOL	C + P	P	S	IPN	SS
SBR [53]	M	FOL	C + F	P	S+Sub	IN	DS + SS
Scallop [58]	P	FOL	M + P	P	S	LpN	DS
SL [135]	M	P	C + P	P	S	LpN	SS
Slash [113]	P+M	FOL	S + P	P	S	LpN	DS +SS
TensorLog [18]	P	R	M + P	P	S	LpN	DS + KGC

[Marra *et al*, AIJ 2024]

Many thanks!

- **Bosch:** Youmna Ismaeil, Tran-Trung Kien, Huu Mai Tan, Tobias Kaminski, Mohamed Gad-Elrab, Coung Xuan Chou, Csaba Domokos, Jan-Hendrik Metzen, Evgeny Kharlamov
 - Bosch domain experts
- **TU Vienna:** Thomas Eiter, Nysret Musliu, Johannes Oetsch, Nelson Higuera, Tobias Geibinger, Medina Andresel
- **KU Leuven:** Hendrik Blockeel
- **MPI:** Gerhard Weikum, Thinh Vinh Ho, Hiba Arnaout
- **Other collaborators:**
 - Nitisha Jain, Peter Skocovski, Pasquale Minervini, Thinh Vinh Ho, Daniel Daza, Michael Cohez, Paul Groth, Zico Kolter, Po-wei Wang



Towards Neuro-Symbolic AI Applications at Bosch: Opportunities and Challenges

Daria Stepanova
Bosch Center for AI

Interested in a master thesis, PhD position or full time opportunities at Bosch, please reach out at daria.stepanova@de.bosch.com

References

- [Ho *et al*, WWW 2022] Vinh Thinh Ho, Daria Stepanova, Dragan Milchevski, Jannik Strötgen, Gerhard Weikum: Enhancing Knowledge Bases with Quantity Facts. WWW 2022: 893-901
- [Wang *et al*, ICLR 2020] Po-Wei Wang, Daria Stepanova, Csaba Domokos, J. Zico Kolter: Differentiable learning of numerical rules in knowledge graphs. ICLR 2020
- [Jain *et al*, ISWC 2021] Nitisha Jain, Trung-Kien Tran, Mohamed H. Gad-Elrab, Daria Stepanova: Improving Knowledge Graph Embeddings with Ontological Reasoning. ISWC 2021: 410-426
- [Andresel *et al*, CIKM 2023] Medina Andresel, Trung-Kien Tran, Csaba Domokos, Pasquale Minervini, Daria Stepanova: Combining Inductive and Deductive Reasoning for Query Answering over Incomplete Knowledge Graphs. CIKM 2023: 15-24
- [Ismaeil *et al*, ISWC 2023] Youmna Ismaeil, Daria Stepanova, Trung-Kien Tran, Hendrik Blockeel: FeaBI: A Feature Selection-Based Framework for Interpreting KG Embeddings. ISWC 2023: 599-617
- [Daza *et al*, TLMR 2026] Daniel Daza, C. Xuan Chu, T.K. Tran, Daria Stepanova, Michael Cochez, Paul Groth: Explaining Graph Neural Networks for Node Similarity on Graphs. TLMR 2026.
- [Chu *et al*, ESWC 2022 posters] Cuong Chu Xuan, Mohamed H. Gad-Elrab, Trung-Kien Tran, Marvin Schiller, Evgeny Kharlamov, Daria Stepanova: Supplier Optimization at Bosch with Knowledge Graphs and Answer Set Programming. ESWC (Satellite Events) 2023: 200-204
- [Qi *et al*, ICLR 2025] Zhenting Qi, Hongyin Luo, Xuliang Huang, Zhuokai Zhao, Yibo Jiang, Xiangjun Fan, Himabindu Lakkaraju, James R. Glass: Quantifying Generalization Complexity for Large Language Models. ICLR 2025
- [Fan *et al*, ACL 2024] Lizhou Fan, Wenyue Hua, Lingyao Li, Haoyang Ling, Yongfeng Zhang: NPHardEval: Dynamic Benchmark on Reasoning Ability of Large Language Models via Complexity Classes. ACL (1) 2024: 4092-4114
- [Chen *et al*, ICLR 2026] Hongzheng Chen, Yingheng Wang, Yaohui Cai, Hins Hu, Jiajie Li, Shirley Huang, Chenhui Deng, Rongjian Liang, Shufeng Kong, Haoxing Ren, Samitha Samaranayake, Carla P. Gomes, Zhiru Zhang: HeuriGym: An Agentic Benchmark for LLM-Crafted Heuristics in Combinatorial Optimization. ICLR 2026
- [Hazra *et al*, Arxiv 2025] Rishi Hazra, Gabriele Venturato, Pedro Zuidberg Dos Martires, Luc De Raedt: Have Large Language Models Learned to Reason? A Characterization via 3-SAT Phase Transition. CoRR abs/2504.03930 (2025)
- [Peng *et al*, COLM 2024] Binghui Peng, Srin Narayanan, Christos H. Papadimitriou: On Limitations of the Transformer Architecture. CoRR abs/2402.08164 (2024)
- [Cheng *et al*, Arxiv 2025] Lijie Chen, Binghui Peng, Hongxun Wu: Theoretical limitations of multi-layer Transformer. CoRR abs/2412.02975 (2024)
- [Thomm *et al*, NeurIPS 2024] Jonathan Thomm, Giacomo Camposampiero, Aleksandar Terzic, Michael Hersche, Bernhard Schölkopf, Abbas Rahimi: Limits of Transformer Language Models on Learning to Compose Algorithms. NeurIPS 2024

References

- [Eiter *et al*, TPLP 2023] Thomas Eiter, Tobias Geibinger, Nysret Musliu, Johannes Oetsch, Peter Skocovský, Daria Stepanova: Answer-Set Programming for Lexicographical Makespan Optimisation in Parallel Machine Scheduling. *Theory Pract. Log. Program.* 23(6): 1281-1306 (2023)
- [Eiter *et al*, AIJ 2024] Thomas Eiter, Tobias Geibinger, Nelson Higuera Ruiz, Nysret Musliu, Johannes Oetsch, Dave Pfliegler, Daria Stepanova: Adaptive large-neighbourhood search for optimisation in answer-set programming. *Artif. Intell.* 337: 104230 (2024)
- [Eiter *et al*, IJCAI 2023] Thomas Eiter, Tobias Geibinger: Explaining Answer-Set Programs with Abstract Constraint Atoms. *IJCAI 2023*: 3193-3202
- [Eiter *et al*, JELIA 2023] Thomas Eiter, Tobias Geibinger, Johannes Oetsch: Contrastive Explanations for Answer-Set Programs. *JELIA 2023*: 73-89
- [Geibinger *et al*, TAASP 2024] Tobias Geibinger, Johannes Oetsch and Tobias Kaminski. Explanations for Guess-and-Check ASP Encodings Using an LLM. *TAASP (Extended Abstract)*. 2024
- [Herud *et al*, ICLP 2022] Konstantin Herud, Joachim Baumeister, Orkunt Sabuncu, Torsten Schaub: Conflict Handling in Product Configuration using Answer Set Programming. *ICLP Workshops 2022*
- [Szeider, Arxiv 2025] Stefan Szeider: MCP-Solver: Integrating Language Models with Constraint Programming Systems. *CoRR abs/2501.00539* (2025)
- [Metzen *et al*, ICCV 2023] Jan Hendrik Metzen, Robin Hutmacher, N. Grace Hua, Valentyn Boreiko, Dan Zhang: Identification of Systematic Errors of Image Classifiers on Rare Subgroups. *ICCV 2023*: 5041-5050
- [Ismaeil *et al*, ISWC 2025] Youmna Ismaeil, Jan-Hendrik Metzen, Tran-Trung Kien, Hendrik Bloeckeel, Daria Stepanova. Beyond Manual Labels: Unsupervised Graph-Based Explanations for Error Analysis in Image Classifiers. *ISWC 2025*
- [Colevatti *et al*, NeSy 2024] Michele Collevati, Thomas Eiter, Nelson Higuera: Leveraging Neurosymbolic AI for Slice Discovery. *NeSy (1)* 2024: 403-418
- [Eiter *et al*, AAAI 2022] Thomas Eiter, Tobias Geibinger, Nelson Higuera Ruiz, Nysret Musliu, Johannes Oetsch, Daria Stepanova: Large-Neighbourhood Search for Optimisation in Answer-Set Solving. *AAAI 2022*: 5616-5625
- [Eiter *et al*, KR 2022] Thomas Eiter, Tobias Geibinger, Nelson Higuera Ruiz, Nysret Musliu, Johannes Oetsch, Daria Stepanova: ALASPO: An Adaptive Large-Neighbourhood ASP Optimiser. *KR 2022*
- [Eiter *et al*, KR 2021] Thomas Eiter, Tobias Geibinger, Nysret Musliu, Johannes Oetsch, Peter Skocovský, Daria Stepanova: Answer-Set Programming for Lexicographical Makespan Optimisation in Parallel Machine Scheduling. *KR 2021*: 280-290

References

- [Gad-Elrab *et al*, ISWC 2020] Mohamed H. Gad-Elrab, Daria Stepanova, Trung-Kien Tran, Heike Adel, Gerhard Weikum: ExCut: Explainable Embedding-Based Clustering over Knowledge Graphs. ISWC (1) 2020: 218-237
- [Ott *et al*, CIKM 2023] Simon Ott, Patrick Betz, Daria Stepanova, Mohamed H. Gad-Elrab, Christian Meilicke, Heiner Stuckenschmidt: Rule-based Knowledge Graph Completion with Canonical Models. CIKM 2023: 1971-1981
- [Ismaeil *et al*, ISWC 2022] Youmna Ismaeil, Daria Stepanova, Trung-Kien Tran, Piyapat Saranrittichai, Csaba Domokos, Hendrik Blockeel: Towards Neural Network Interpretability Using Commonsense Knowledge Graphs. ISWC 2022: 74-90
- [Zou *et al*, 2026] Huanghaohe Zou, Peng Han, Emad Nazerian, Alex Q. Huang: PCBSchemaGen: Constraint-Guided Schematic Design via LLM for Printed Circuit Boards (PCB). CoRR abs/2602.00510 (2026)