Neuro-Symbolic AI: Explainability and Deep Deductive Reasoning

Pascal Hitzler
Data Semantics Laboratory (DaSe Lab)
Kansas State University

http://www.daselab.org

Neuro-symbolic AI

Publications on neuro-symbolic AI in major conferences (research papers only):

conference	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	total
ICML	0	0	0	0	0	1	3	2	5	6	17
NeurIPS	0	0	0	0	0	0	0	4	2	4	10
AAAI	0	0	0	0	0	1	0	1	1	1	4
IJCAI	1	0	0	0	0	0	2	2	0	2	7
ICLR	N/A	N/A	0	0	0	0	1	1	1	3	6
total	1	0	0	0	0	2	6	10	9	16	44

See

Md Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, Pascal Hitzler

Neuro-Symbolic Artificial Integration: Current Trends

AI Communications, to appear; https://arxiv.org/abs/2105.05330 for more analysis.

Neural

- **Refers to computational abstractions of (natural) neural network systems.**
- **Prominently includes Artificial Neural Networks and Deep Learning as machine learning paradigms.**
- **More generally sometimes referred to as *connectionist systems*.**

- **Prominent applications come from the machine learning world.**

- **And of course, there is the current deep learning hype.**

Symbolic

- Refers to (computational) symbol manipulations of all kind.

- Graphs and trees, traversal, data structure operations.
- Knowledge representation in explicit symbolic form (data base, ontology, knowledge graph)
- Inductive and statistical inference.
- Formal logical (deductive or abductive) reasoning.

- Prominent applications all over computer science, including expert systems (and their modern versions), information systems, data management, added value of data annotation, etc.

- Semantic Web data is inherently symbolic.

Neuro-Symbolic

Computer Science perspective:

- **Let's try to get the best of both worlds:**
  - **very powerful machine learning paradigm**
  - **robust to data noise**
  - **easy to understand and assess by humans**
  - **good at symbol manipulation**
  - **work seamlessly with background (domain) knowledge**

- **How to do that?**
  - **Endow connectionist systems with symbolic components?**
  - **Add connectionist learning to symbolic reasoners?**

Some Background

Workshop Series on Neural-Symbolic Learning and Reasoning, since 2005.
Joint with Artur d'Avila Garcez.
http://neural-symbolic.org/

Barbara Hammer and Pascal Hitzler (eds), Perspectives of Neural-Symbolic Integration, Springer, 2007

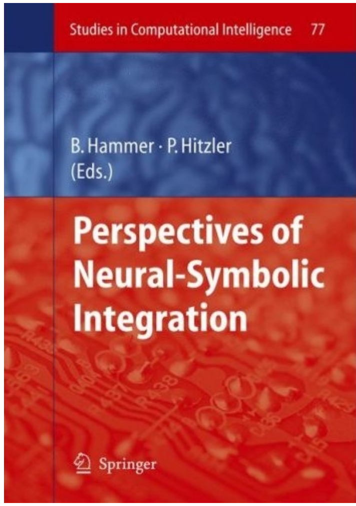

Neural-Symbolic Learning and Reasoning: A Survey and Interpretation

Tarek R. Besold, Artur d'Avila Garcez, Sebastian Bader,
Howard Bowman, Pedro Domingos, Pascal Hitzler,
Kai-Uwe Kuehnberger, Luis C. Lamb, Daniel Lowd,
Priscila Machado Vieira Lima, Leo de Penning, Gadi Pinkas,
Hoifung Poon, Gerson Zaverucha

https://arxiv.org/abs/1711.03902 (2017)

Ilaria Tiddi, Freddy Lecue, Pascal Hitzler (eds.), Knowledge Graphs for eXplainable Artificial Intelligence: Foundations, Applications and Challenges. Studies on the Semantic Web Vol. 47, IOS Press, 2020.

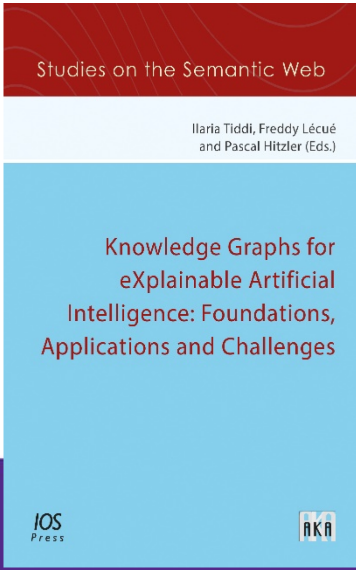

2022 Book (just out!)

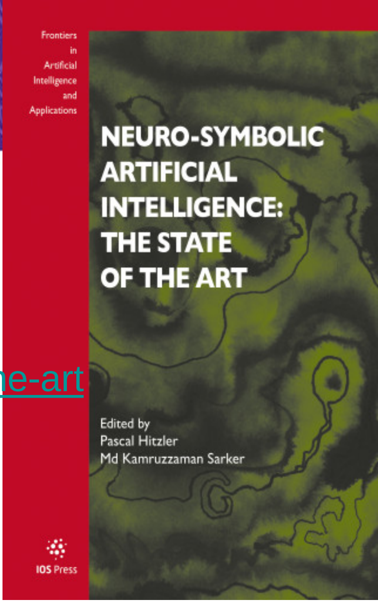

Neuro-symbolic Artificial Intelligence: The State of the Art

Pascal Hitzler and Md Kamruzzaman Sarker, editors

Fontriers in AI and Applications Vol. 342, IOS Press, Amsterdam, 2022

https://www.iospress.com/catalog/books/neuro-symbolic-artificial-intelligence-the-state-of-the-art

Deep Deductive Reasoners

Monireh Ebrahimi, Aaron Eberhart, Federico Bianchi, Pascal Hitzler,
Towards Bridging the Neuro-Symbolic Gap: Deep Deductive Reasoners.
Applied Intelligence 51 (9), 6326-6348, 2021.

Pascal Hitzler, Frank van Harmelen
A reasonable Semantic Web.
Semantic Web 1 (1-2), 39-44, 2010.

Deep Deductive Reasoners

- We trained deep learning systems to do deductive reasoning.

- Why is this interesting?
  - For dealing with noisy data (where symbolic reasoners do very poorly).
  - For speed, as symbolic algorithms are of very high complexity.
  - Out of principle because we want to learn about the capabilities of deep learning for complicated cognitive tasks.

  - To perhaps begin to understand how our (neural) brains can learn to do highly symbolic tasks like formal logical reasoning, or in more generality, mathematics.
    A fundamental quest in Cognitive Science.

Reasoning as Classification

- Given a set of logical formulas (a theory).
- Any formula expressible over the same language is either
  - a logical consequence or
  - not a logical consequence.
- This can be understood as a classification problem for machine learning.
- It turns out to be a really hard machine learning problem.

Knowledge Materialization

- Given a set of logical formulas (a theory).
- Produce all logical consequences under certain constraints.
- Without the qualifier this is in general not possible as the set of all logical consequences is infinite.
- So we have to constrain to consequences of, e.g., a certain syntactic form. For relatively simple logics, this is often reasonably possible.

Published deep deductive reasoning work

paper	logic	transfer	generative	scale	performance
[12]	RDFS	yes	no	moderate	high
[25]	RDFS	no	yes	low	high
[10]	$\mathcal{EL}^+$	no	yes	moderate	low
[20]	OWL RL	no*	no	low	high
[6]	FOL	no	yes	very low	high
(new)	RDFS	yes	yes	moderate	high
(new)	EL+	yes	yes	moderate	high

[12]: Ebrahimi, Sarker, Bianchi, Xie, Eberhart, Doran, Kim, Hitzler, AAAI-MAKE 2021

[25]: Makni, Hendler, SWJ 2019

[10]: Eberhart, Ebrahimi, Zhou, Shimizu, Hitzler, AAAI-MAKE 2020

[20]: Hohenecker, Lukasiewicz, JAIR 2020

[6]: Bianchi, Hitzler, AAAI-MAKE 2019

(new): Ebrahimi, Eberhart, Hitzler, June 2021

Knowledge Graphs and Ontologies

Pascal Hitzler, Semantic Web: A Review of the Field.
Communications of the ACM 64 (2), 76-82, 2021.

Knowledge Graphs and Ontologies (Schemas)

Knowledge Graphs (and their schemas) are made to enable easier

- **data sharing**
- **data discovery**
- **data integration**
- **data reuse**

Google Knowledge Graph

Laura Kelly

Governor of Kansas

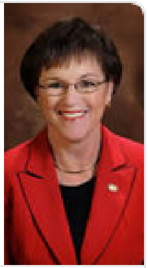

Laura Kelly is an American politician serving as the 48th governor of Kansas since 2019. A member of the Democratic Party, she represented the 18th district in the Kansas Senate from 2005 to 2019. Kelly ran for governor in the 2018 election and defeated the Republican nominee, Kansas Secretary of State Kris Kobach. Wikipedia

Born: January 24, 1950 (age 69 years), New York, NY

Spouse: Ted Daughety

Party: Democratic Party

Office: Governor of Kansas since 2019

Education: Indiana University, Bradley University, Indiana University Bloomington

Children: Kathleen Daughety, Molly Daughety

→ hasEducation →

Indiana University

iu.edu

Indiana University is a multi-campus public university system in the state of Indiana, United States. Indiana University has a combined student body of more than 110,000 students, which includes approximately 46,000 students enrolled at the Indiana University Bloomington campus. Wikipedia

Mascot: Referred to as "The Hoosiers"

Endowment: 1.986 billion USD

Students: 110,436 university-wide

President: Michael McRobbie

Academic staff: 8,733 university-wide

Subsidiaries: Indiana University Bloomington, MORE

→ hasPresident →

Michael McRobbie

President of Indiana University

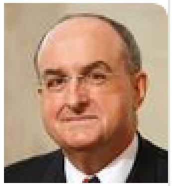

president.iu.edu

Michael Alexander McRobbie AO is an Australian-American computer scientist, educator and academic administrator. He became the eighteenth president of Indiana University on July 1, 2007. Wikipedia

Born: October 11, 1950 (age 69 years), Melbourne, Australia

Spouse: Laurie Burns (m. 2005)

Education: The Australian National University, The University of Queensland

Books: Automated Theorem-proving in Non-classical Logics, Automated Deduction - Cade-13

Knowledge Graphs

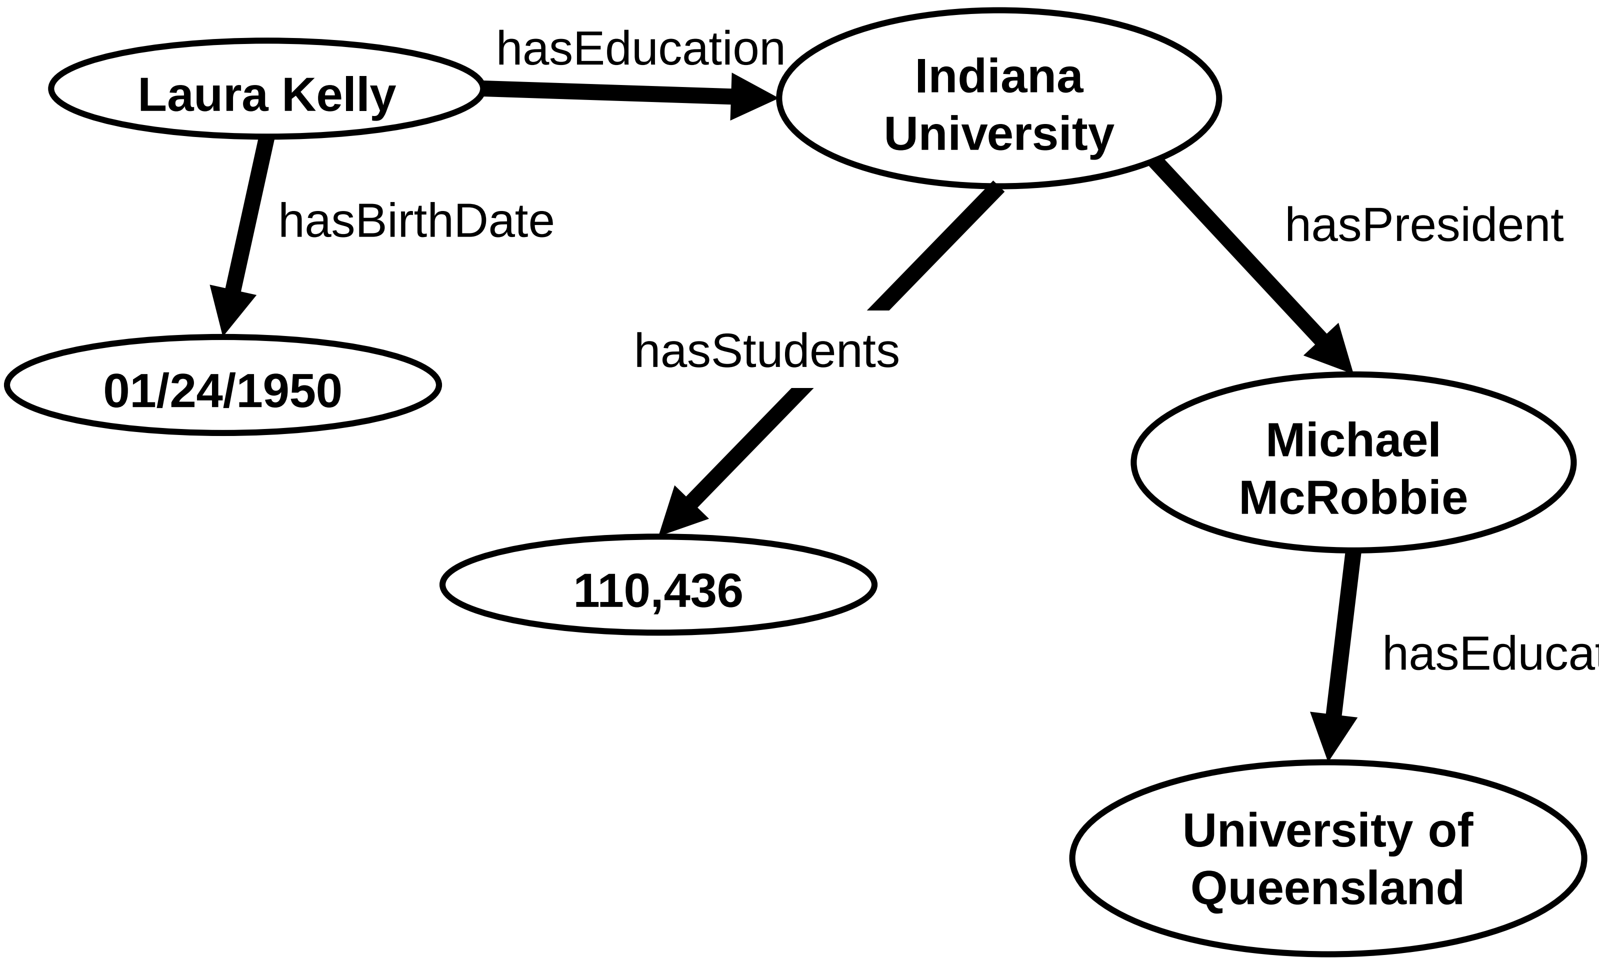

Schema (as diagram)

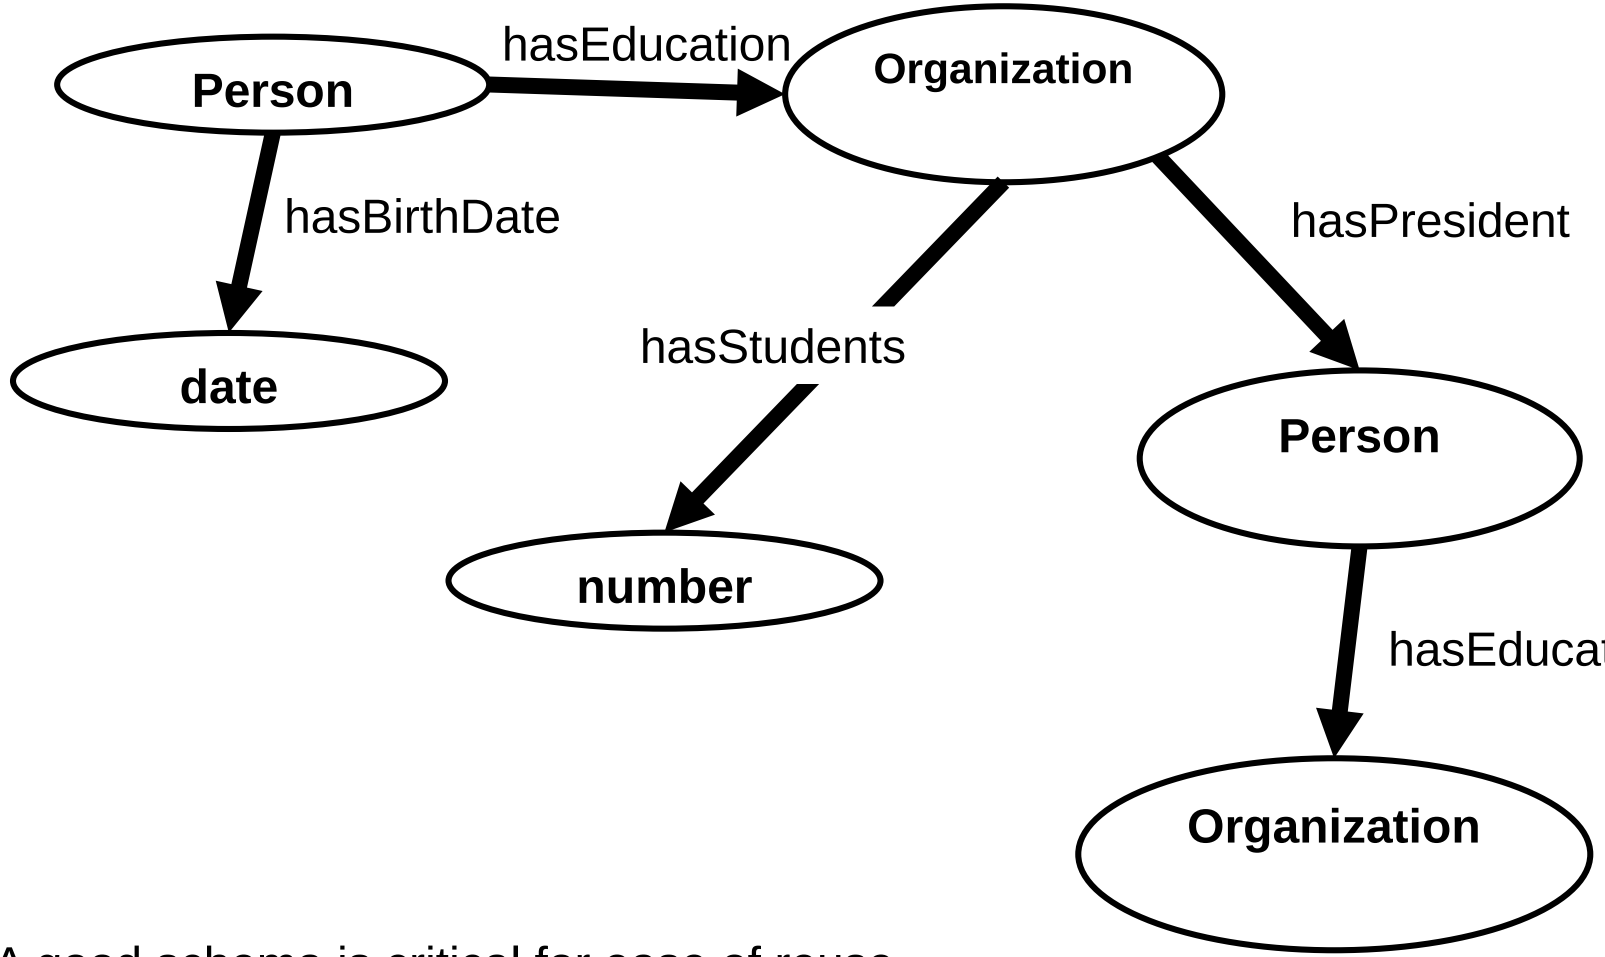

A good schema is critical for ease of reuse

W3C Standards

W3C Recommendation

RDF 1.1 Concepts and Abstract Syntax

W3C Recommendation 25 February 2014

This version:
http://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/
Latest published version:
http://www.w3.org/TR/rdf11-concepts/
Previous version:
http://www.w3.org/TR/2014/PR-rdf11-concepts-20140109/
Previous Recommendation:
http://www.w3.org/TR/rdf-concepts
Editors:
Richard Cyganiak, DERI, NUI Galway
David Wood, 3 Round Stones
Markus Lanthaler, Graz University of Technology

DaSe Lab

Both established 2004
as versions 1.0.

W3C Recommendation

OWL 2 Web Ontology Language Primer (Second Edition)

W3C Recommendation 11 December 2012

This version:
http://www.w3.org/TR/2012/REC-owl2-primer-20121211/
Latest version (series 2):
http://www.w3.org/TR/owl2-primer/
Latest Recommendation:
http://www.w3.org/TR/owl-primer
Previous version:
http://www.w3.org/TR/2012/PER-owl2-primer-20121018/
Editors:
Pascal Hitzler, Wright State University
Markus Krötzsch, University of Oxford
Bijan Parsia, University of Manchester
Peter F. Patel-Schneider, Nuance Communications
Sebastian Rudolph, FZI Research Center for Information

PRACTICE

Industry-Scale Knowledge Graphs: Lessons and Challenges

By Natasha Noy, Yuqing Gao, Anshu Jain, Anant Narayanan, Alan Patterson, Jamie Taylor
Communications of the ACM, August 2019, Vol. 62 No. 8, Pages 36-43
10.1145/3331166
Comments

VIEW AS: SHARE:

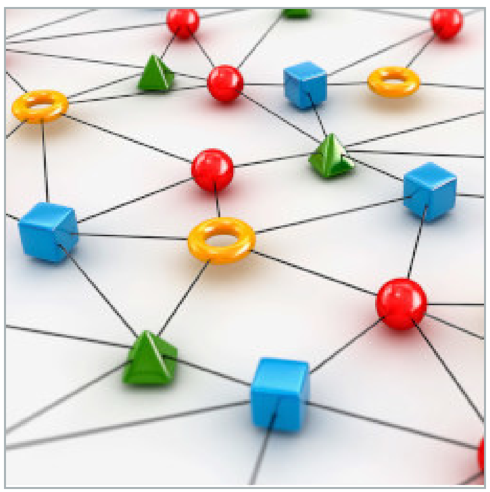

Credit: Adempercem / Stutterstock

Knowledge graphs are critical to many enterprises today: They provide the structured data and factual knowledge that drive many products and make them more intelligent and "magical."

In general, a knowledge graph describes objects of interest and connections between them. For example, a knowledge graph may have nodes for a movie, the actors in this movie, the director, and so on. Each node may have properties such as an actor's name and age. There may be nodes for multiple movies involving a particular actor. The user can then traverse the knowledge graph to collect information on all the movies in which the actor appeared or, if applicable, directed.

Many practical implementations impose constraints on the links in knowledge graphs by defining a *schema* or *ontology*. For example, a link from a movie to its director must connect an object of type Movie to an object of type Person. In some cases the links themselves might have their own properties: a link connecting an actor and a movie might have the name of the specific role the actor

ARTICLE CONTENTS:

MORE NEWS & OPINIONS
MIT Robot Could Help People

Back to:

Deep Deductive Reasoners

Monireh Ebrahimi, Aaron Eberhart, Federico Bianchi, Pascal Hitzler,
Towards Bridging the Neuro-Symbolic Gap: Deep Deductive Reasoners.
Applied Intelligence 51 (9), 6326-6348, 2021.

Pascal Hitzler, Frank van Harmelen
A reasonable Semantic Web.
Semantic Web 1 (1-2), 39-44, 2010.

Published deep deductive reasoning work

paper	logic	transfer	generative	scale	performance
[12]	RDFS	yes	no	moderate	high
[25]	RDFS	no	yes	low	high
[10]	$\mathcal{EL}^+$	no	yes	moderate	low
[20]	OWL RL	no*	no	low	high
[6]	FOL	no	yes	very low	high
(new)	RDFS	yes	yes	moderate	high
(new)	EL+	yes	yes	moderate	high

[12]: Ebrahimi, Sarker, Bianchi, Xie, Eberhart, Doran, Kim, Hitzler, AAAI-MAKE 2021

[25]: Makni, Hendler, SWJ 2019

[10]: Eberhart, Ebrahimi, Zhou, Shimizu, Hitzler, AAAI-MAKE 2020

[20]: Hohenecker, Lukasiewicz, JAIR 2020

[6]: Bianchi, Hitzler, AAAI-MAKE 2019

(new): Ebrahimi, Eberhart, Hitzler, June 2021

RDFS Reasoning using Memory Networks

Monireh Ebrahimi, Md Kamruzzaman Sarker, Federico Bianchi, Ning Xie,
Aaron Eberhart, Derek Doran, Hyeongsik Kim, Pascal Hitzler,
Neuro-Symbolic Deductive Reasoning for Cross-Knowledge Graph Entailment.
In: Proc. AAAI-MAKE 2021.

additional analysis by Sulogna Chowdhury, Aaron Eberhart
and Brayden Pankaskie

RDF reasoning

- [Note: RDF is one of the simplest useful knowledge representation languages that is not propositional.]

- Think knowledge graph.
- Think node-edge-node triples such as

BarackObama	rdf:type	President
BarackObama	husbandOf	MichelleObama
President	rdfs:subClassOf	Human
husbandOf	rdfs:subPropertyOf	spouseOf

- Then there is a (fixed, small) set of inference rules, such as
  rdf:type(x,y) AND rdfs:subClassOf(y,z)THEN rdf:type(x,z)

RDF reasoning

- Essentially, RDF reasoning is Datalog reasoning restricted to:
  - Unary and binary predicates only.
  - A fixed set of rules that are not facts.
- You can try the following:
  - Use a vector embedding for one RDF graph.
  - Create all logical consequences.
  - Throw n% of them away.
  - Use the rest to train a DL system.
  - Check how many of those you threw away can be recovered this way.

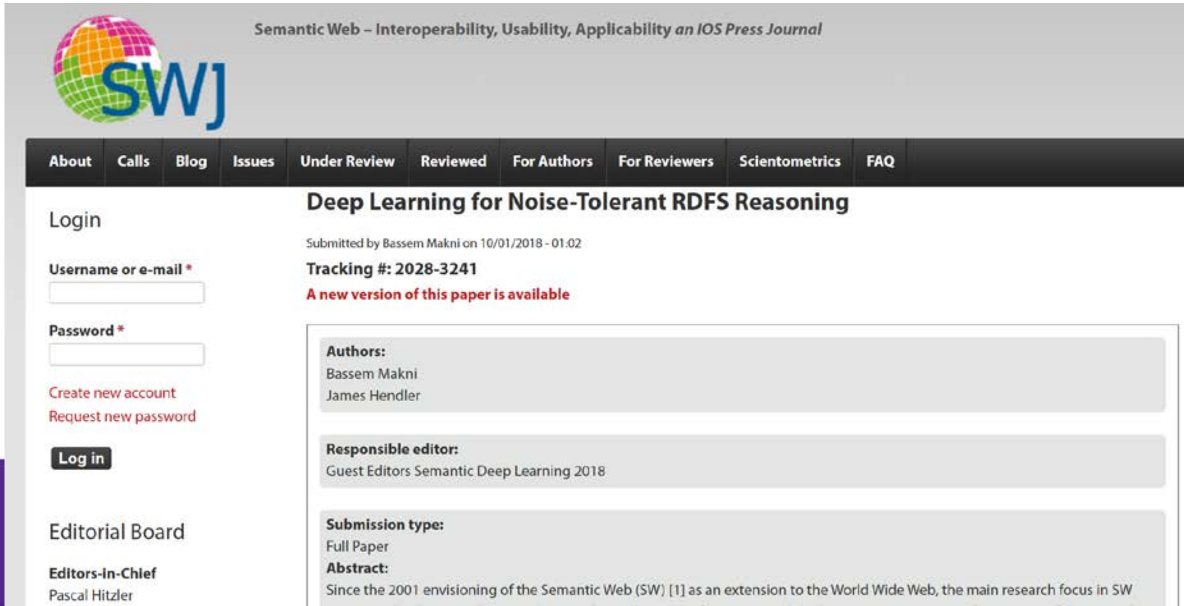

RDF reasoning

- The problem with the approach just described:
  – It works only with that graph.

- What you'd really like to do is:
  – Train a deep learning system so that you can present a new, unseen graph to it, and it can correctly derive the deductively inferred triples.

- Note:
  – You don't know the IRIs in the graph up front. The only overlap may or may not be the IRIs in the rdf/s namespace.
  – You don't know up front how "deep" the reasoning needs to be.
  – There is no lack of training data, it can be auto-generated.

Representation

- Goal is to be able to reason over unseen knowledge graphs. I.e. the out-of-vocabulary problem needs addressing.

- Normalization of vocabulary (i.e., it becomes shared vocabulary across all input knowledge graphs.
- One vocabulary item becomes a one-hot vector (dimension d, number of normalized vocabulary terms)
- One triple becomes a 3 x d matrix.
- The knowledge graph becomes an n x 3 x d tensor (n is the number of knowledge graph triples)

- Knowledge graph is stored in “memory”

Mechanics

- An attention mechanism retrieves memory slots useful for finding the correct answer to a query.
- These are combined with the query and run through a (learned) matrix to retrieve a new (processed) query.

- This is repeated (in our experiment with 10 “hops”).
- The final out put is a yes/no answer to the query.

Memory Network based on MemN2N

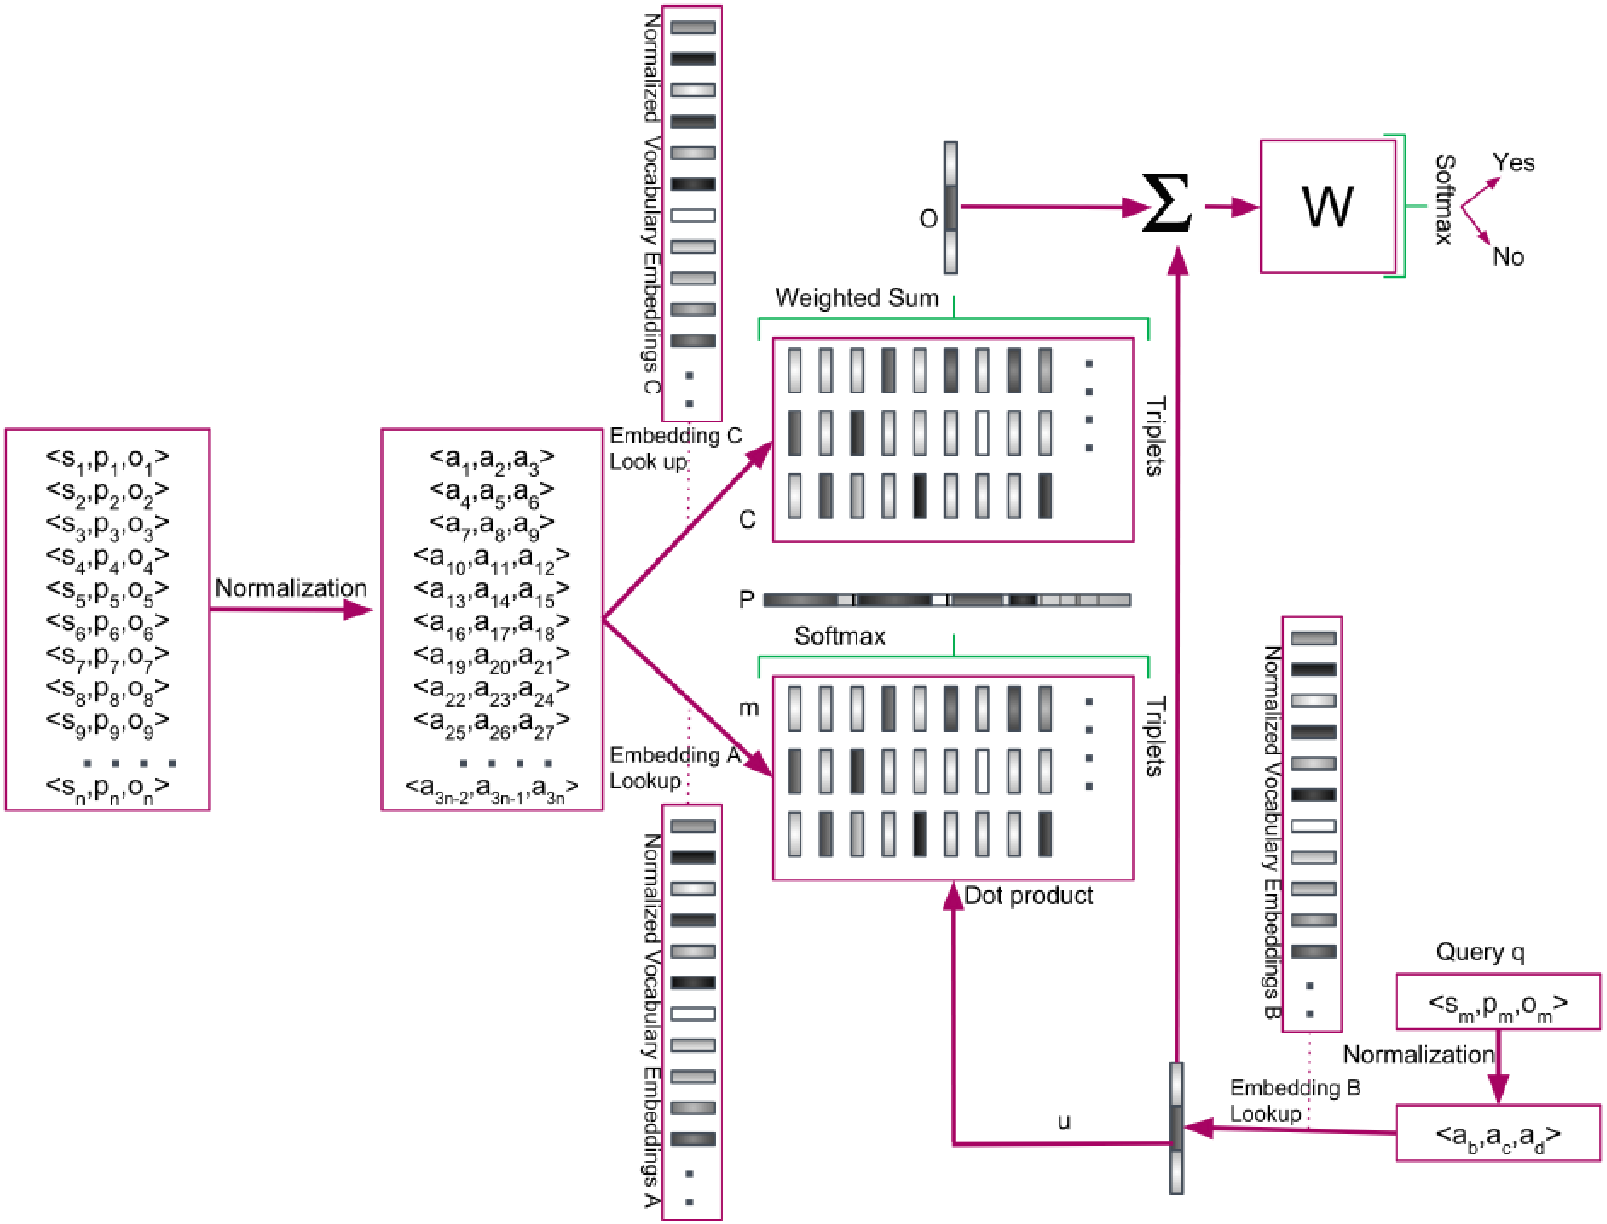

Experiments: Performance

Test Dataset	#KG	Base						Inferred						Invalid
		#Facts	#Ent.	%Class	%Indv	%R.	%Axiom.	#Facts	#Ent.	%Class	%Indv	%R.	%Axiom.	#Facts
OWL-Centric	2464	996	832	14	19	3	0	494	832	14	0.01	1	20	462
Linked Data	20527	999	787	3	22	5	0	124	787	3	0.006	1	85	124
OWL-Centric Test Set	21	622	400	36	41	3	0	837	400	36	3	1	12	476
Synthetic Data	2	752	506	52	0	1	0	126356	506	52	0	1	0.07	700

Table 2: Statistics of various datasets used in experiments

Baseline: non-normalized embeddings, same architecture

Training Dataset	Test Dataset	Valid Triples Class			Invalid Triples Class			Accuracy
		Precision	Recall /Sensitivity	F-measure	Precision	Recall /Specificity	F-measure	
OWL-Centric Dataset	Linked Data	93	98	96	98	93	95	**96**
OWL-Centric Dataset (90%)	OWL-Centric Dataset (10%)	88	91	89	90	88	89	**90**
OWL-Centric Dataset	OWL-Centric Test Set [b]	79	62	68	70	84	76	**69**
OWL-Centric Dataset	Synthetic Data	65	49	40	52	54	42	**52**
OWL-Centric Dataset	Linked Data [a]	54	98	70	91	16	27	86
OWL-Centric Dataset [a]	Linked Data [a]	62	72	67	67	56	61	91
OWL-Centric Dataset(90%) [a]	OWL-Centric Dataset(10%) [a]	79	72	75	74	81	77	80
OWL-Centric Dataset	OWL-Centric Test Set [ab]	58	68	62	62	50	54	58
OWL-Centric Dataset [a]	OWL-Centric Test Set [ab]	77	57	65	66	82	73	73
OWL-Centric Dataset	Synthetic Data [a]	70	51	40	47	52	38	51
OWL-Centric Dataset [a]	Synthetic Data [a]	67	23	25	52	80	62	50
Baseline								
OWL-Centric Dataset	Linked Data	73	98	83	94	46	61	43
OWL-Centric Dataset (90%)	OWL-Centric Dataset (10%)	84	83	84	84	84	84	82
OWL-Centric Dataset	OWL-Centric Test Set [b]	62	84	70	80	40	48	61
OWL-Centric Dataset	Synthetic Data	35	41	32	48	55	45	48

[a] More Tricky Nos & Balanced Dataset
[b] Completely Different Domain.

Table 3: Experimental results of proposed model

Experiments: Reasoning Depth

DaSe Lab

Test Dataset	Hop 0			Hop 1			Hop 2			Hop 3			Hop 4		
	P	R	F	P	R	F	P	R	F	P	R	F	P	R	F
Linked Data[a]	0	0	0	80	99	88	89	97	93	77	98	86	-	-	-
Linked Data[b]	2	0	0	82	91	86	89	98	93	79	100	88	-	-	-
OWL-Centric[c]	19	5	9	31	75	42	78	80	78	48	47	44	4	34	6
Synthetic	32	46	33	31	87	38	66	55	44	25	45	32	29	46	33

Test Dataset	Hop 5			Hop 6			Hop 7			Hop 8			Hop 9			Hop 10		
	P	R	F	P	R	F	P	R	F	P	R	F	P	R	F	P	R	F
Linked Data[a]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Linked Data[b]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
OWL-Centric[c]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Synthetic	26	46	33	25	46	33	25	46	33	24	43	31	25	43	31	22	36	28

[a] LemonUby Ontology
[b] Agrovoc Ontology
[c] Completely Different Domain

Table 4: Experimental results over each reasoning hop

Dataset	Hop 1	Hop 2	Hop 3	Hop 4	Hop 5	Hop 6	Hop 7	Hop 8	Hop 9	Hop 10
OWL-Centric [a]	8%	67%	24%	0.01%	0%	0%	0%	0%	0%	0%
Linked Data[b]	31%	50%	19%	0%	0%	0%	0%	0%	0%	0%
Linked Data[c]	34%	46%	20%	0%	0%	0%	0%	0%	0%	0%
OWL-Centric[d]	5%	64%	30%	1%	0%	0%	0%	0%	0%	0%
Synthetic Data	0.03%	1.42%	1%	1.56%	3.09%	6.03%	11.46%	20.48%	31.25%	23.65%

[a] *Training Set*
[b] LemonUby Ontology
[c] Agrovoc Ontology
[d] Completely Different Domain

Table 5: Data distribution per knowledge graph over each reasoning hop

Training time: just over a full day

Generative RDFS Reasoning using Pointer Networks

Monireh Ebrahimi, Aaron Eberhart, Pascal Hitzler
On the Capabilities of Pointer Networks for Deep Deductive Reasoning
https://arxiv.org/abs/2106.09225

Pointer Networks

- **Pointer Networks 'point' to input elements!**
- **Ptr-Net approach specifically targets problems whose outputs are discrete and correspond to positions in the input.**
- **At each time step, the distribution of the attention is the answer!**
- **Application:**
  - **NP-hard Travelling Salesman Problem (TSP)**
  - **Delaunay Triangulation**
  - **Convex Hull**
  - **Text Summarization**
  - **Code completion**
  - **Dependency Parsing**

Pointer Networks for Reasoning

- **To mimic human reasoning behaviour where one can learn to choose a set of symbols in different locations and copy these symbols to suitable locations to generate new logical consequences based on a set of predefined logical entailment rules**

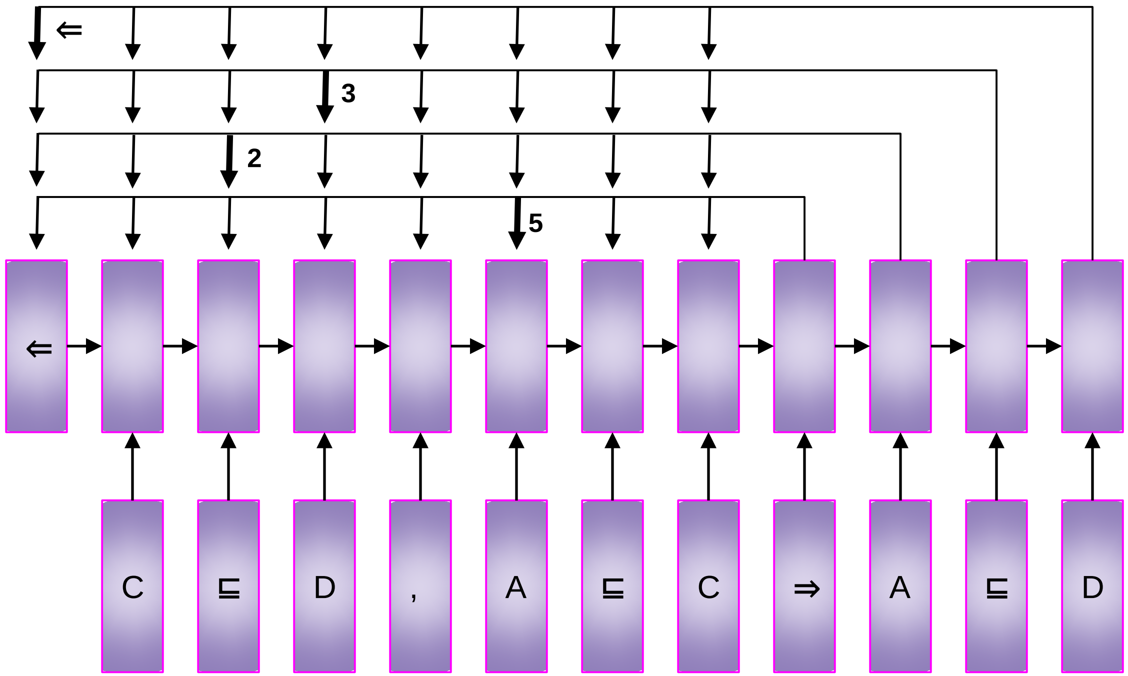

$$C \sqsubseteq D, A \sqsubseteq C \mapsto A \sqsubseteq D$$

Results without transfer

Logic	KG Size	Pointer Networks		Transformer			LSTM
		SubWordText	Tokenizer	Normalized	Not-Normalized		
					SubWordText	Tokenizer	
RDF	3 - 735	87%	99%	5%	25%	4%	0.17%

- On RDF, slightly outperforms [Hendler Makni SWJ 2019] approach.
- Our approach is a more straightforward application.
- Evaluation is on the same dataset.

Results with transfer

Table 6 Exact Match Accuracy Results for Transfer Learning/Representation: SubWord-Text Tokenization Encoding

Train \ Test	LUBM	Awards	University
LUBM	*	75%	78%
Awards	79%	*	77%
University	81%	82%	*

Table 7 Exact Match Accuracy Results for Transfer Learning/ Representation: Whitespace Tokenization Encoding

Train \ Test	LUBM	Awards	University
LUBM	*	61%	47%
Awards	96%	*	84%
University	82%	88%	*

Completion Reasoning Emulation for the Description Logic EL+

Aaron Eberhart, Monireh Ebrahimi, Lu Zhou, Cogan Shimizu, Pascal Hitzler, Completion Reasoning Emulation for the Description Logic EL+. In: Andreas Martin, Knut Hinkelmann, Hans-Georg Fill, Aurona Gerber, Doug Lenat, Reinhard Stolle, Frank van Harmelen (eds.), Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice, AAAI-MAKE 2020, Palo Alto, CA, USA, March 23-25, 2020, Volume I.

EL+ is essentially OWL 2 EL

$CX \sqsubseteq CY$
$CX \sqcap CY \sqsubseteq CZ$
$CX \sqsubseteq \exists RY.CZ$
$\exists RX.CY \sqsubseteq CZ$
$RX \sqsubseteq RY$
$RX \circ RY \sqsubseteq RZ$

Table 2: $\mathcal{EL}^+$ Completion Rules

(1)	$A \sqsubseteq C$	$C \sqsubseteq D$		$\models A \sqsubseteq D$
(2)	$A \sqsubseteq C_1$	$A \sqsubseteq C_2$	$C_1 \sqcap C_2 \sqsubseteq D$	$\models A \sqsubseteq D$
(3)	$A \sqsubseteq C$	$C \sqsubseteq \exists R.D$		$\models A \sqsubseteq \exists R.D$
(4)	$A \sqsubseteq \exists R.B$	$B \sqsubseteq C$	$\exists R.C \sqsubseteq D$	$\models A \sqsubseteq D$
(5)	$A \sqsubseteq \exists S.D$	$S \sqsubseteq R$		$\models A \sqsubseteq \exists R.D$
(6)	$A \sqsubseteq \exists R_1.C$	$C \sqsubseteq \exists R_2.D$	$R_1 \circ R_2 \sqsubseteq R$	$\models A \sqsubseteq \exists R.D$

Table 1: $\mathcal{EL}^+$ Semantics

Description	Expression	Semantics
Individual	a	$a \in \Delta^{\mathcal{I}}$
Top	$\top$	$\Delta^{\mathcal{I}}$
Bottom	$\bot$	$\emptyset$
Concept	C	$C^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$
Role	R	$R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$
Conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
Existential Restriction	$\exists R.C$	$\{ a \mid$ there is $b \in \Delta^{\mathcal{I}}$ such that $(a, b) \in R^{\mathcal{I}}$ and $b \in C^{\mathcal{I}} \}$
Concept Subsumption	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
Role Subsumption	$R \sqsubseteq S$	$R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$
Role Chain	$R_1 \circ \cdots \circ R_n \sqsubseteq R$	$R_1^{\mathcal{I}} \circ \cdots \circ R_n^{\mathcal{I}} \subseteq R^{\mathcal{I}}$

with $\circ$ signifying standard binary composition

Results

Table 7: Average Precision Recall and F1-score For each Distance Evaluation

	Atomic Levenshtein Distance			Character Levenshtein Distance			Predicate Distance		
	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score
	Synthetic Data								
Piecewise Prediction	0.138663	0.142208	0.140412	0.138663	0.142208	0.140412	0.138646	0.141923	0.140264
Deep Prediction	**0.154398**	**0.156056**	**0.155222**	**0.154398**	**0.156056**	**0.155222**	**0.154258**	**0.155736**	**0.154993**
Flat Prediction	0.140410	0.142976	0.141681	0.140410	0.142976	0.141681	0.140375	0.142687	0.141521
Random Prediction	0.010951	0.0200518	0.014166	0.006833	0.012401	0.008811	0.004352	0.007908	0.007908
	SNOMED Data								
Piecewise Prediction	0.010530	0.013554	0.011845	0.010530	0.013554	0.011845	0.010521	0.013554	0.011839
Deep Prediction	**0.015983**	0.0172811	**0.016595**	**0.015983**	0.017281	**0.016595**	**0.015614**	0.017281	**0.016396**
Flat Prediction	0.014414	**0.018300**	0.016112	0.0144140	**0.018300**	0.016112	0.013495	**0.018300**	0.015525
Random Prediction	0.002807	0.006803	0.003975	0.001433	0.003444	0.002023	0.001769	0.004281	0.002504

Generative EL Reasoning using Pointer Networks

Monireh Ebrahimi, Aaron Eberhart, Pascal Hitzler
On the Capabilities of Pointer Networks for Deep Deductive Reasoning
https://arxiv.org/abs/2106.09225

Results with transfer

Logic	KG Size	Pointer Networks		Transformer			LSTM
		SubWordText	Tokenizer	Normalized	Not-Normalized		
					SubWordText	Tokenizer	
[redacted]							
ER	40	73%	**73%**	8%	8%	0.4 %	0%
	50	68%	**68%**	11%	11%	0.3%	0%
	120	49%	**49%**	15%	NA	NA	0%

- **same architecture as before**

Explaining Deep Learning via Symbolic Background Knowledge

Md. Kamruzzaman Sarker, Ning Xie, Derek Doran, Michael Raymer, Pascal Hitzler, Explaining Trained Neural Networks with Semantic Web Technologies: First Steps. In: Tarek R. Besold, Artur S. d'Avila Garcez, Isaac Noble (eds.), Proceedings of the Twelfth International Workshop on Neural-Symbolic Learning and Reasoning, NeSy 2017, London, UK, July 17-18, 2017. CEUR Workshop Proceedings 2003, CEUR-WS.org 2017

Md Kamruzzaman Sarker, Pascal Hitzler, Efficient Concept Induction for Description Logics. In: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 – February 1, 2019. AAAI Press 2019 , pp. 3036-3043.

Md Kamruzzaman Sarker, Joshua Schwartz, Pascal Hitzler, Lu Zhou, Srikanth Nadella, Brandon Minnery, Ion Juvina, Michael L. Raymer, William R. Aue, Wikipedia Knowledge Graph for Explainable AI. In: Boris Villazón-Terrazas, Fernando Ortiz-Rodríguez, Sanju M. Tiwari, Shishir K. Shandilya (eds.), Knowledge Graphs and Semantic Web. Second Iberoamerican Conference and First Indo-American Conference, KGSWC 2020, Mérida, Mexico, November 26-27, 2020, Proceedings. Communications in Computer and Information Science, vol. 1232, Springer, Heidelberg, 2020, pp. 72-87.

Explainable AI

- Explain behavior of trained (deep) NNs.

- Idea:
  - Use background knowledge in the form of linked data and ontologies to help explain.
  - Link inputs and outputs to background knowledge.
  - Use a symbolic learning system to generate an explanatory theory.

- We have key components for this now, but it's still early stages.

Concept

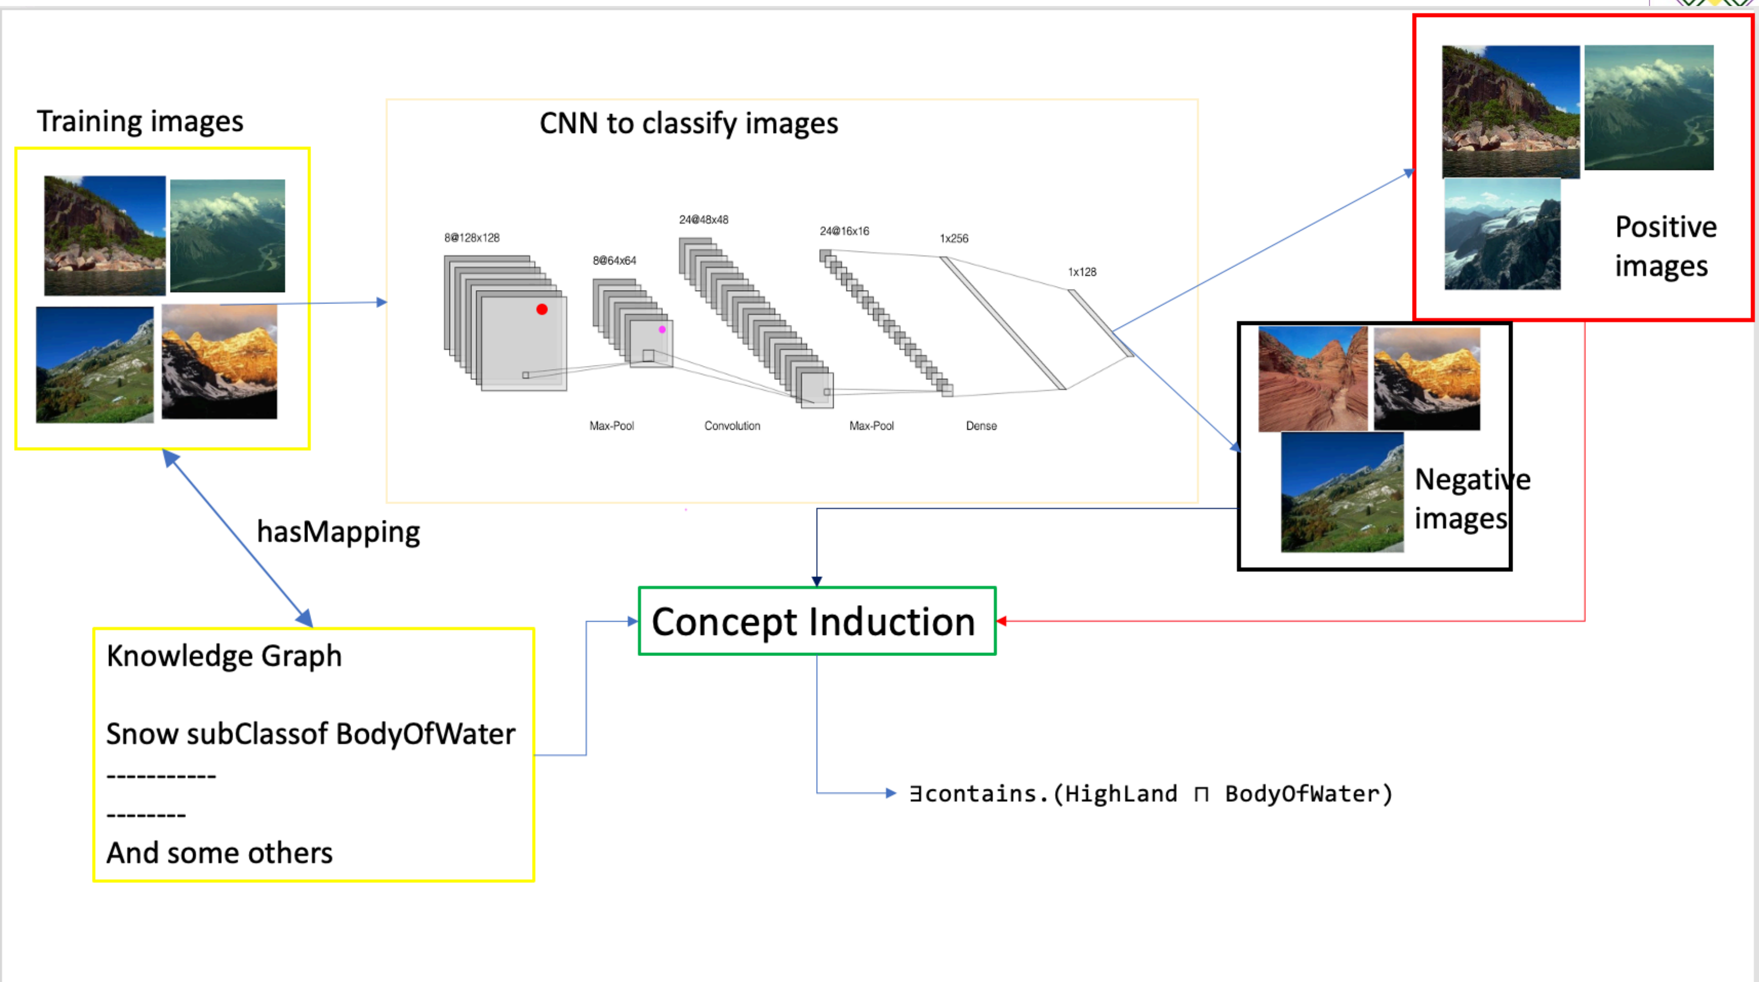

Training images
CNN to classify images
8@128x128
8@64x64
24@48x48
24@16x16
1x256
1x128
Max-Pool
Convolution
Max-Pool
Dense
Positive images
Negative images
hasMapping
Concept Induction
Knowledge Graph
Snow subClassof BodyOfWater

And some others
∃contains.(HighLand ⊓ BodyOfWater)

DL-Learner [Lehmann, Hitzler]

Approach similar to inductive logic programming, but using Description Logics (the logic underlying OWL).

Positive examples: **negative examples:**

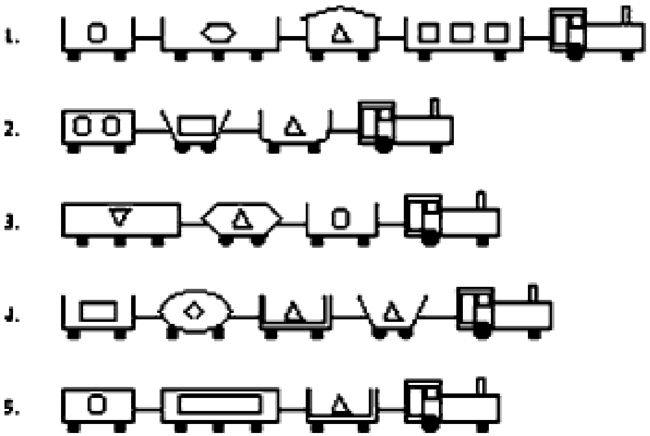

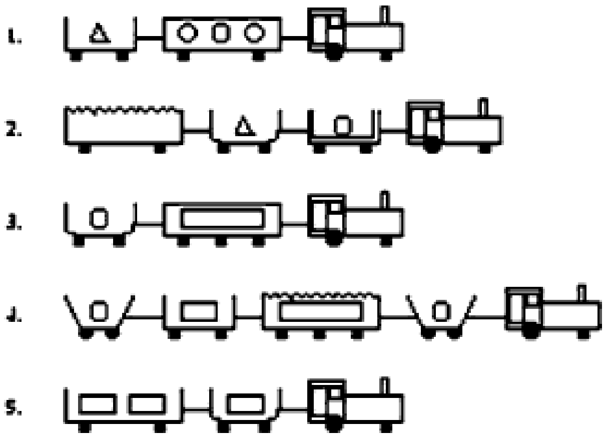

Task: find a class description (logical formula) which separates positive and negative examples.

DL-Learner

Positive examples:

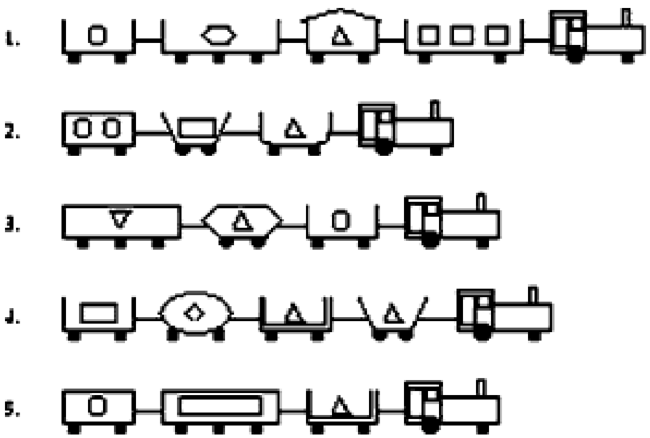

negative examples:

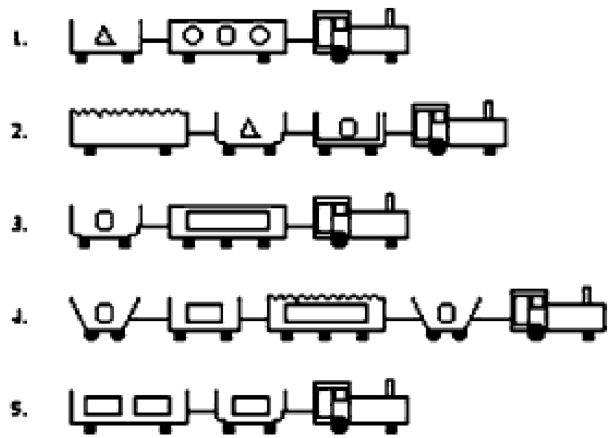

DL-Learner result: $\exists \text{hasCar}.(\text{Closed} \sqcap \text{Short})$

In FOL:

$$\{x \mid \exists y(\text{hasCar}(x, y) \wedge \text{Closed}(y) \wedge \text{Short}(y))\}$$

Scalability Issues with DL-Learner

- For large-scale experiments, DL-Learner took 2 hours or more for one run.
- We knew we needed at least thousands of runs.

- So we needed a more scalable solution.
- The provably correct algorithms have very high complexity.

- Hence we had to develop a heuristic which trades (some) correctness for speed.

- It is also currently restricted to using a class hierarchy as underlying knowledge base.

ECII algorithm and system

- **We thus implemented our own system, ECII (Efficient Concept Induction from Instances) which trades some correctness for speed. [Sarker, Hitzler, AAAI-19]**

DaSe Lab

Experiment Name	Number of Logical Axioms	Runtime (sec)					Accuracy (α_3)		Accuracy α_2			
		DL[a]	DL FIC(1)[b]	DL FIC(2)[c]	ECII DF[d]	ECII KCT[e]	DL[a]	ECII DF[d]	DL FIC(1)[b]	DL FIC(2)[c]	ECII DF[d]	ECII KCT[e]
Yinyang_examples	157	0.065	0.0131	0.019	0.089	0.143	1.000	0.610	1.000	1.000	0.799	1.000
Trains	273	0.01	0.020	0.047	0.05	0.095	1.000	1.000	1.000	1.000	1.000	1.000
Forte	341	2.5	1.169	6.145	0.95	0.331	0.965	0.642	0.875	0.875	0.733	1.000
Poker	1,368	0.066	0.714	0.817	1	0.281	1.000	1.000	0.981	0.984	1.000	1.000
Moral Reasoner	4,666	0.1	3.106	4.154	5.47	6.873	1.000	0.785	1.000	1.000	1.000	1.000
ADE20k I	4,714	577.3[f]	4.268	31.887	1.966	23.775	0.926	0.416	0.263	0.814	0.744	1.000
ADE20k II	7,300	983.4[f]	16.187	307.65	20.8	293.44	1.000	0.673	0.413	0.413	0.846	0.900
ADE20k III	12,193	4,500[g]	13.202	263.217	51	238.8	0.375	0.937	0.375	0.375	0.930	0.937
ADE20k IV	47,468	4,500[g]	93.658	523.673	116	423.349	0.375	NA	0.608	0.608	0.660	0.608

[a] DL : DL-Learner
[b] DL FIC (1) : DL-Learner fast instance check with runtime capped at execution time of ECII DF
[c] DL FIC (2) : DL-Learner fast instance check with runtime capped at execution time of ECII KCT
[d] ECII DF : ECII default parameters
[e] ECII KCT : ECII keep common types and other default parameters
[f] Runtimes for DL-Learner were capped at 600 seconds.
[g] Runtimes for DL-Learner were capped at 4,500 seconds.

ECII vs. DL-Learner

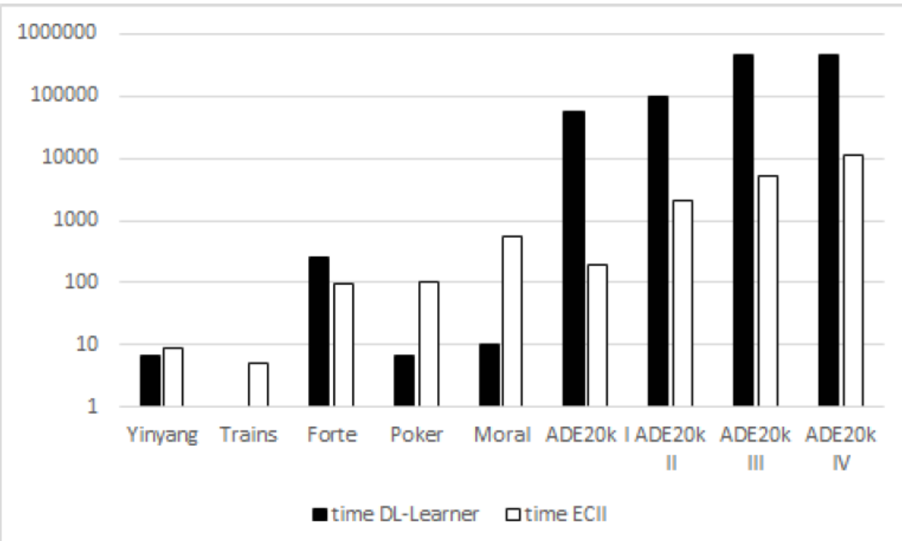

Figure 1: Runtime comparison between DL-Learner and ECII. The vertical scale is logarithmic in hundredths of seconds, and note that DL-Learner runtime has been capped at 4,500 seconds for ADE20k III and IV. For ADE20k I it was capped at each run at 600 seconds.

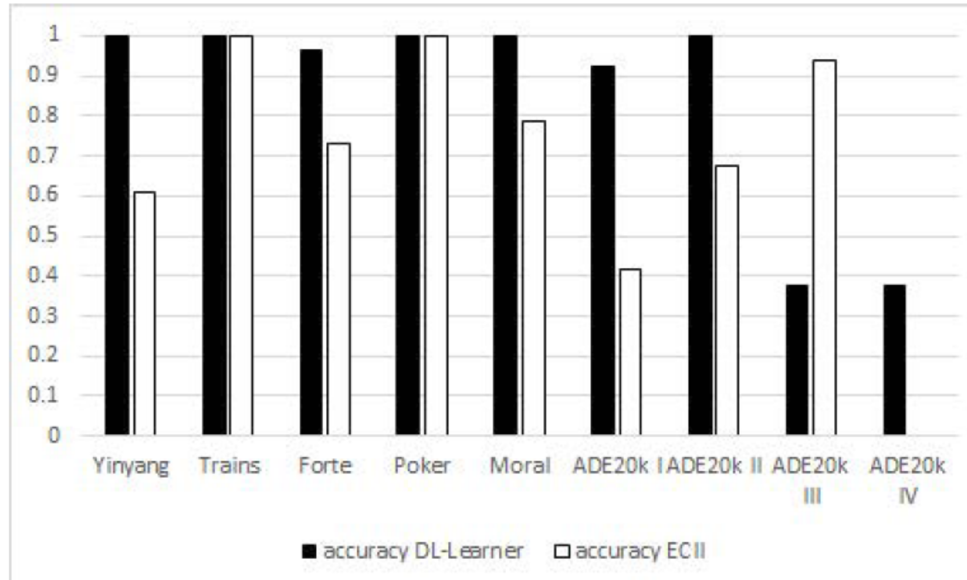

Figure 2: Accuracy (α_3) comparison between DL-Learner and ECII. For ADE20k IV it was not possible to compute an accuracy score within 3 hours for ECII as the input ontology was too large.

Proof of Concept Experiment

Positive:

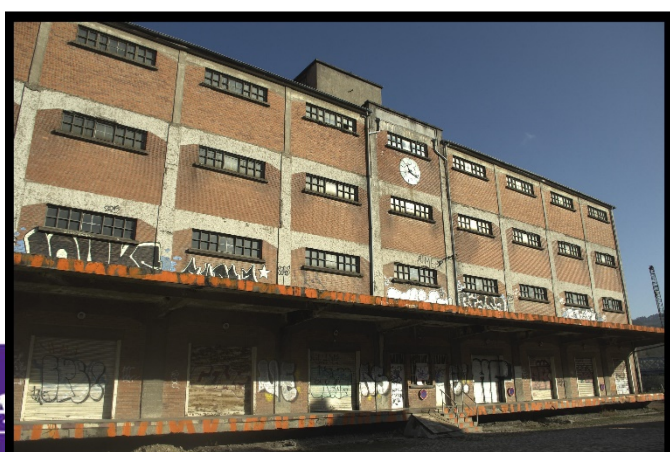

Negative:

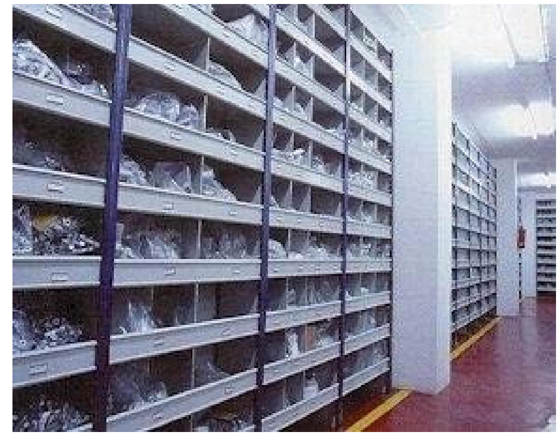

Images

Come from the MIT ADE20k dataset
http://groups.csail.mit.edu/vision/datasets/ADE20K/
They come with annotations of objects in the picture:

```
001 # 0 # 0 # sky # sky # ""
002 # 0 # 0 # road, route # road # ""
005 # 0 # 0 # sidewalk, pavement # sidewalk # ""
006 # 0 # 0 # building, edifice # building # ""
007 # 0 # 0 # truck, motortruck # truck # ""
008 # 0 # 0 # hovel, hut, hutch, shack, shanty # hut # ""
009 # 0 # 0 # pallet # pallet # ""
011 # 0 # 0 # box # boxes # ""
001 # 1 # 0 # door # door # ""
002 # 1 # 0 # window # window # ""
009 # 1 # 0 # wheel # wheel # ""
```

Mapping to SUMO

Simple approach: for each known object in image, create an individual for the ontology which is in the appropriate SUMO class:

contains road1

contains window1

contains door1

contains wheel1

contains sidewalk1

contains truck1

contains box1

contains building1

SUMO

- Suggested Merged Upper Ontology
  http://www.adampease.org/OP/
- Approx. 25,000 common terms
  covering a wide range of domains
- Centrally, a relatively naïve class hierarchy.
- Objects in image annotations became individuals (constants), which were then typed using SUMO classes.

DL-Learner input

Positive:

img1: road, window, door, wheel, sidewalk, truck, box, building

img2: tree, road, window, timber, building, lumber

img3: hand, sidewalk, clock, steps, door, face, building, window, road

Negative:

img4: shelf, ceiling, floor

img5: box, floor, wall, ceiling, product

img6: ceiling, wall, shelf, floor, product

DL-Learner results include:

$$\exists \text{contains}.\text{Transitway}$$

$$\exists \text{contains}.\text{LandArea}$$

Proof of Concept Experiment

Positive:

Negative:

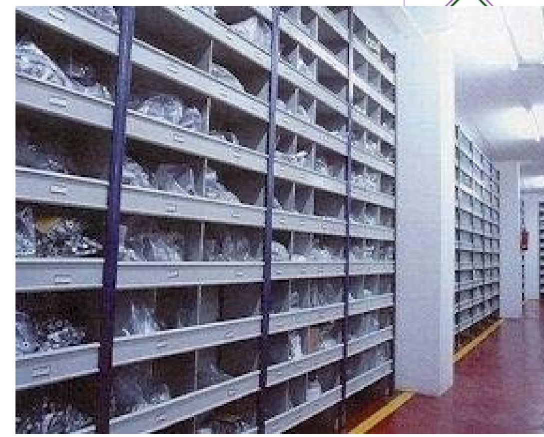

$$\exists \text{contains.Transitway}$$

$$\exists \text{contains.LandArea}$$

Experiment 2

Positive (selection):

$\exists$contains.SentientAgent

Negative (selection):

Experiment 5

Positive:




Negative (selection):



$\exists \text{contains}.\text{BodyOfWater}$

Idea Recap

- Generate explanation of the whole model
- Global explanation

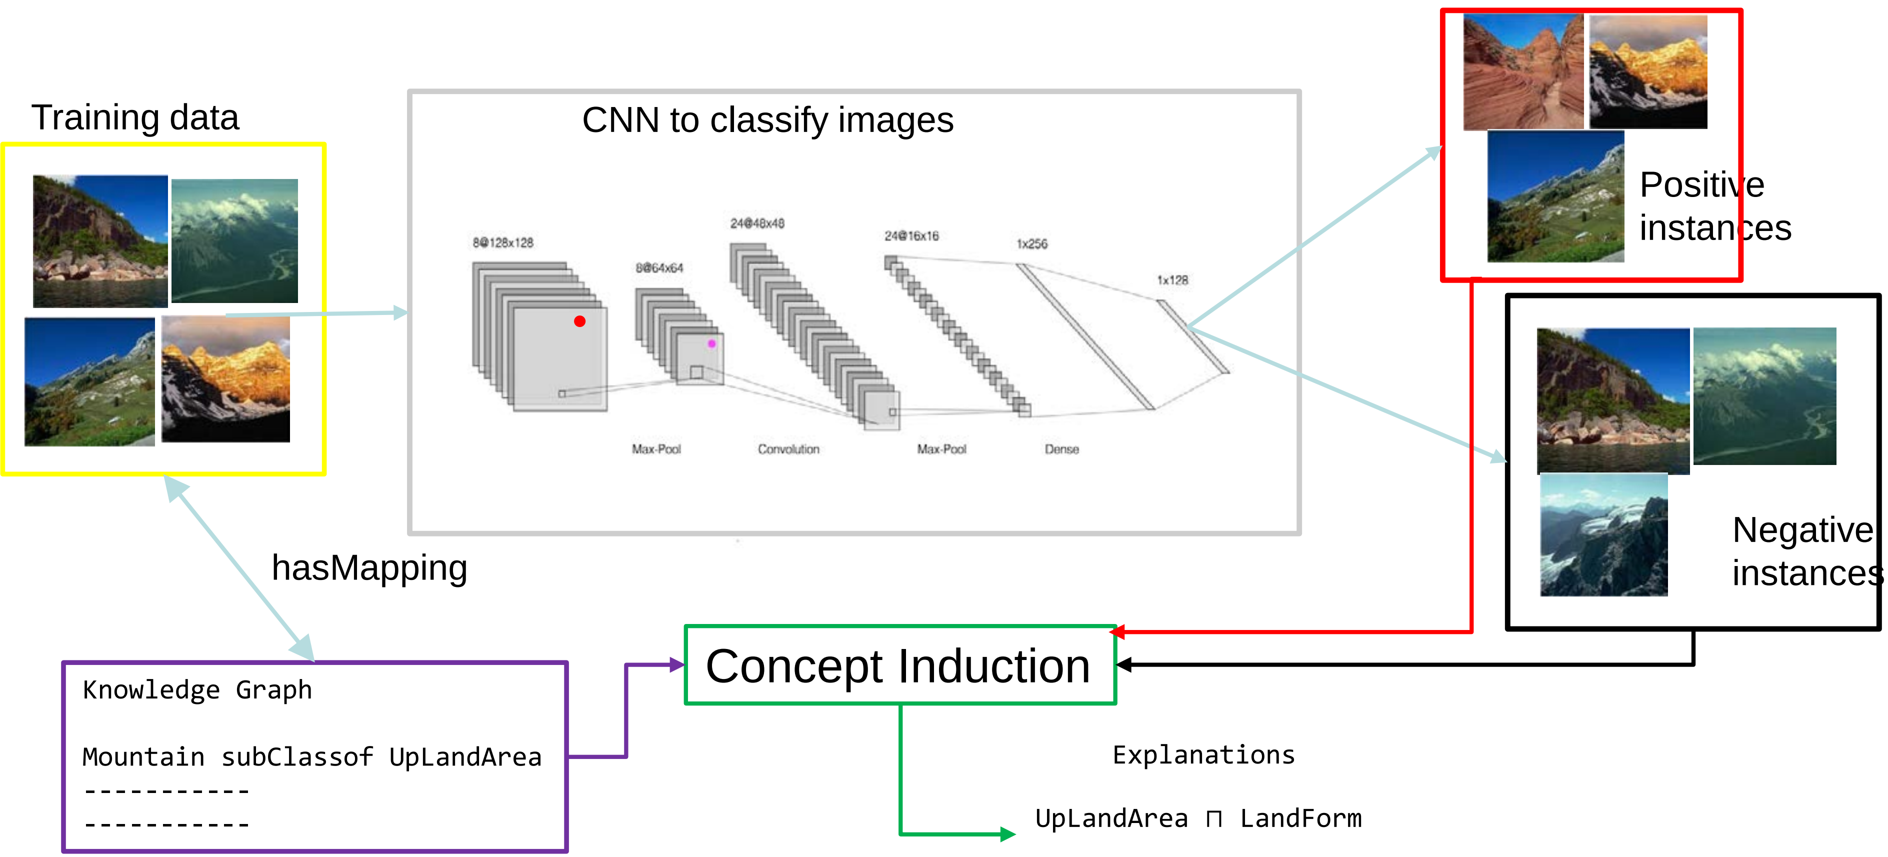

From SUMO to Wikipedia Concept Hierarchy

- Wikipedia CH (curated) produces better coverage score
- Reason behind this is the large number of concepts it has.
  - approx. 2M concepts

Experiment name	#Images	#Positive images	Wikipedia		SUMO	
			#Solution	Coverage	#Solution	Coverage
Market vs. WorkRoom and wareHouse	96	37	286	.72	240	.72
Mountain vs. Market and workRoom	181	85	195	.61	190	.53
OutdoorWarehouse vs. IndoorWarehouse	55	3	128	.94	102	.89
Warehouse vs. Workroom	59	55	268	.56	84	.24
Workroom vs. Warehouse	59	4	128	.93	93	.84

Work in Progress

- Value of Explanations (end-to-end) to
  - humans
  - detect bias
  - improve deep learning accuracy

  - background knowledge challenges

- Explaining hidden neuron activation patterns
  - scalability challenges
  - background knowledge challenges

Conclusions

Conclusions

- Bridging the neuro-symbolic gap is still a major quest.

- Research on Deep Deductive Reasoning is at the heart of neuro-symbolic Artificial Intelligence
  - Research is needed to push the envelope with respect to core aspects such as
  - more complex logics
  - higher reasoning accuracy
  - better transfer
  - scalability
- Knowledge Graphs bear the promise for explainable AI, with explanations from background knowledge.
  - but this is still very much work in progress.

Thanks!

References

Pascal Hitzler, Md Kamruzzaman Sarker (eds.), Neuro-Symbolic Artificial Intelligence – The State of the Art. Frontiers in Artificial Intelligence and Applications Vol. 342, IOS Press, Amsterdam, 2022.

Monireh Ebrahimi, Aaron Eberhart, Federico Bianchi, Pascal Hitzler, Towards Bridging the Neuro-Symbolic Gap: Deep Deductive Reasoners. Applied Intelligence 51 (9), 6326-6348, 2021.

Barbara Hammer and Pascal Hitzler (eds), Perspectives on Neural-Symbolic Integration. Springer, 2007

Tarek R. Besold, Artur d'Avila Garcez, Sebastian Bader, Howard Bowman, Pedro Domingos, Pascal Hitzler, Kai-Uwe Kuehnberger, Luis C. Lamb, Daniel Lowd, Priscila Machado Vieira Lima, Leo de Penning, Gadi Pinkas, Hoifung Poon, Gerson Zaverucha, Neural-Symbolic Learning and Reasoning: A Survey and Interpretation. https://arxiv.org/abs/1711.03902 (2017)

Md Kamruzzaman Sarker, Lu Zhou, Aaron Eberhart, Pascal Hitzler
Neuro-Symbolic Artificial Integration: Current Trends
AI Communications, to appear.

References

Federico Bianchi, Pascal Hitzler, On the Capabilities of Logic Tensor Networks for Deductive Reasoning. In: Andreas Martin, Knut Hinkelmann, Aurona Gerber, Doug Lenat, Frank van Harmelen, Peter Clark (eds.), Proceedings of the AAAI 2019 Spring Symposium on Combining Machine Learning with Knowledge Engineering (AAAI-MAKE 2019) Stanford University, Palo Alto, California, USA, March 25-27, 2019, Stanford University, Palo Alto, California, USA, March 25-27, 2019. CEUR Workshop Proceedings 2350, CEUR-WS.org 2019.

Aaron Eberhart, Monireh Ebrahimi, Lu Zhou, Cogan Shimizu, Pascal Hitzler, Completion Reasoning Emulation for the Description Logic EL+. In: Andreas Martin, Knut Hinkelmann, Hans-Georg Fill, Aurona Gerber, Doug Lenat, Reinhard Stolle, Frank van Harmelen (eds.), Proceedings of the AAAI 2020 Spring Symposium on Combining Machine Learning and Knowledge Engineering in Practice, AAAI-MAKE 2020, Palo Alto, CA, USA, March 23-25, 2020, Volume I.

References

Monireh Ebrahimi, Md Kamruzzaman Sarker, Federico Bianchi, Ning Xie, Aaron Eberhart, Derek Doran, Hyeongsik Kim, Pascal Hitzler, Neuro-Symbolic Deductive Reasoning for Cross-Knowledge Graph Entailment. In: Proc. AAAI-MAKE 2021.

Bassem Makni, James Hendler, Deep learning for noise-tolerant RDFS reasoning. Semantic Web 10(5): 823-862 (2019)

Monireh Ebrahimi, Aaron Eberhart, Pascal Hitzler, On the Capabilities of Pointer Networks for Deep Deductive Reasoning. June 2021. https://arxiv.org/abs/2106.09225

Pascal Hitzler, Frank van Harmelen, A reasonable Semantic Web. Semantic Web 1 (1-2), 39-44, 2010.

Pascal Hitzler, Federico Bianchi, Monireh Ebrahimi, Md Kamruzzaman Sarker, Neural-Symbolic Integration and the Semantic Web. Semantic Web 11 (1), 2020, 3-11.

References

Federico Bianchi, Matteo Palmonari, Pascal Hitzler, Luciano Serafini, Complementing Logical Reasoning with Sub-symbolic Commonsense. In: Paul Fodor, Marco Montali, Diego Calvanese, Dumitru Roman, Rules and Reasoning - Third International Joint Conference, RuleML+RR 2019, Bolzano, Italy, September 16-19, 2019, Proceedings. Lecture Notes in Computer Science 11784, Springer 2019, pp. 161-170.

Sebastian Bader, Pascal Hitzler, Dimensions of neural-symbolic integration – a structured survey. In: S. Artemov, H. Barringer, A. S. d'Avila Garcez, L. C. Lamb and J. Woods (eds). We Will Show Them: Essays in Honour of Dov Gabbay, Volume 1. International Federation for Computational Logic, College Publications, 2005, pp. 167-194.

Pascal Hitzler, Semantic Web: A Review of the Field. Communications of the ACM 64 (2), 76-82, 2021.

Thanks!