Neurosymbolic Artificial Intelligence – some results regarding knowledge graphs

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

http://www.daselab.org

Contents

- Two current trends:
  - Neuro-Symbolic Artificial Intelligence
  - Knowledge Graphs
- And their convergence:
  - Added Value for Deep Learning
    - Example: Explainable AI
  - Added Value for Knowledge Graphs
    - Example: Deep Deductive Reasoning

Neuro-Symbolic Artificial Intelligence

2022 Book

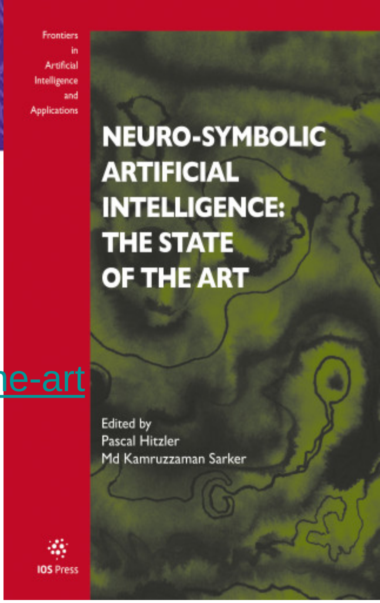

Neuro-symbolic Artificial Intelligence: The State of the Art

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

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)
- 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 / knowledge graphs are 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**

- **But how to do this best?**

Contents

- Two current trends:
  - Neuro-Symbolic Artificial Intelligence
  - Knowledge Graphs
- And their convergence:
  - Added Value for Deep Learning
    - Example: Explainable AI
  - Added Value for Knowledge Graphs
    - Example: Deep Deductive Reasoning

Knowledge Graphs

Google Knowledge Graph

Laura Kelly

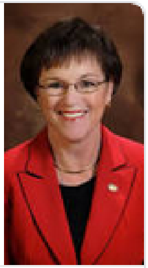

Governor of Kansas

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

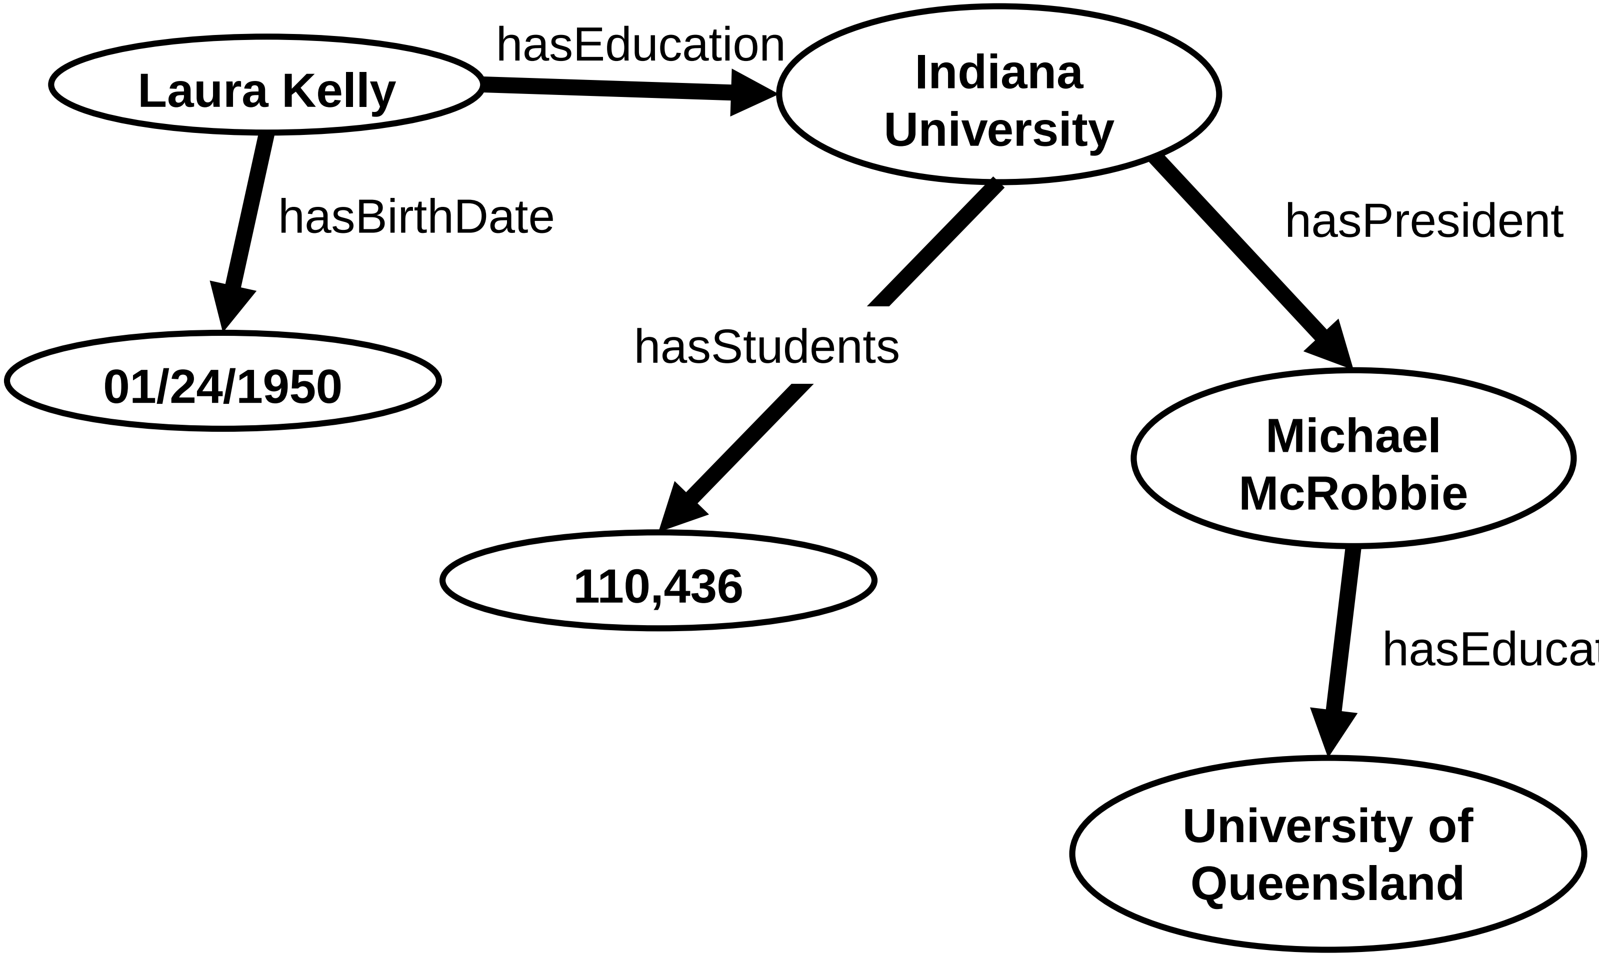

Schema (as diagram), aka Ontology

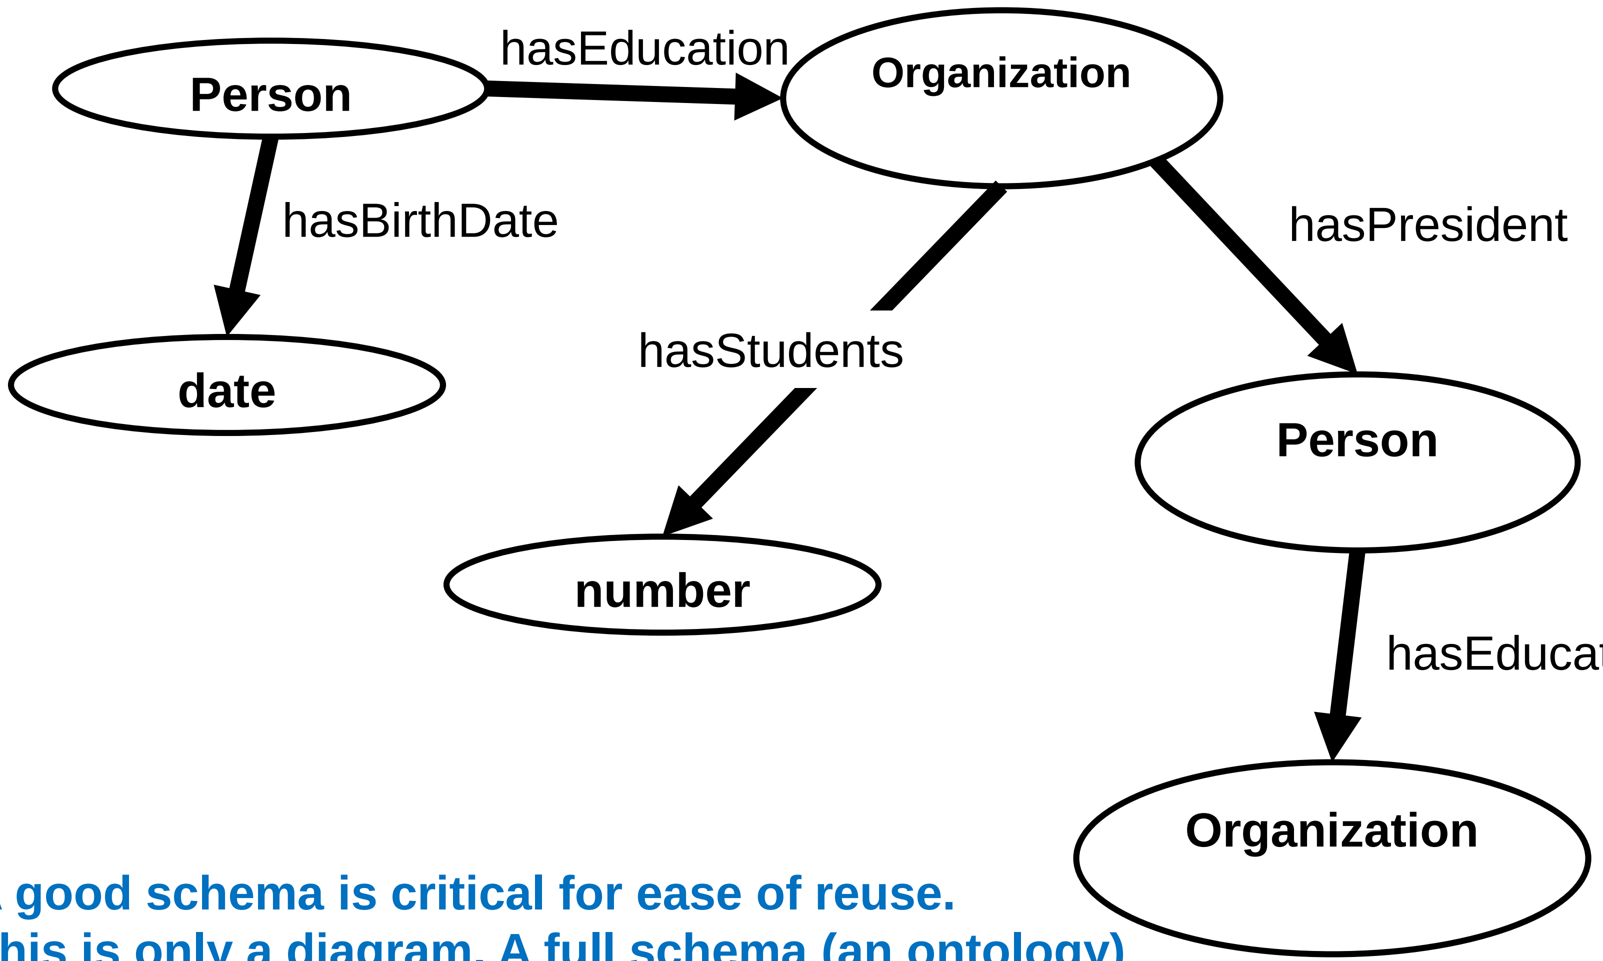

A good schema is critical for ease of reuse.
This is only a diagram. A full schema (an ontology) consists of axioms in a formal logic.

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

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

ARTICLE CONTENTS:

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

MORE NEWS & OPINIONS

MIT Robot Could Help People

collaborative

open

free

multilingu

Welcome to Wikidata

the free knowledge base with 97,501,043 data items that anyone can edit.

Introduction • Project Chat • Community Portal • Help

Want to help translate? Translate the missing messages.

Welcome!

Wikidata is a free and open knowledge base that can be read and edited by both humans and machines.

Wikidata acts as central storage for the **structured data** of its Wikimedia sister projects including Wikipedia, Wikivoyage, Wiktionary, Wikisource, and others.

Wikidata also provides support to many other sites and services beyond just Wikimedia projects! The content of Wikidata is available under a free license, exported using standard formats, and can be interlinked to other open data sets on the linked data web.

Learn about data

New to the wonderful world of data? Develop and improve your data literacy through content designed to get you up to speed and feeling comfortable with the fundamentals in no time.

Item: *Earth (Q2)*

Property: *highest point*

Gartner, 2021

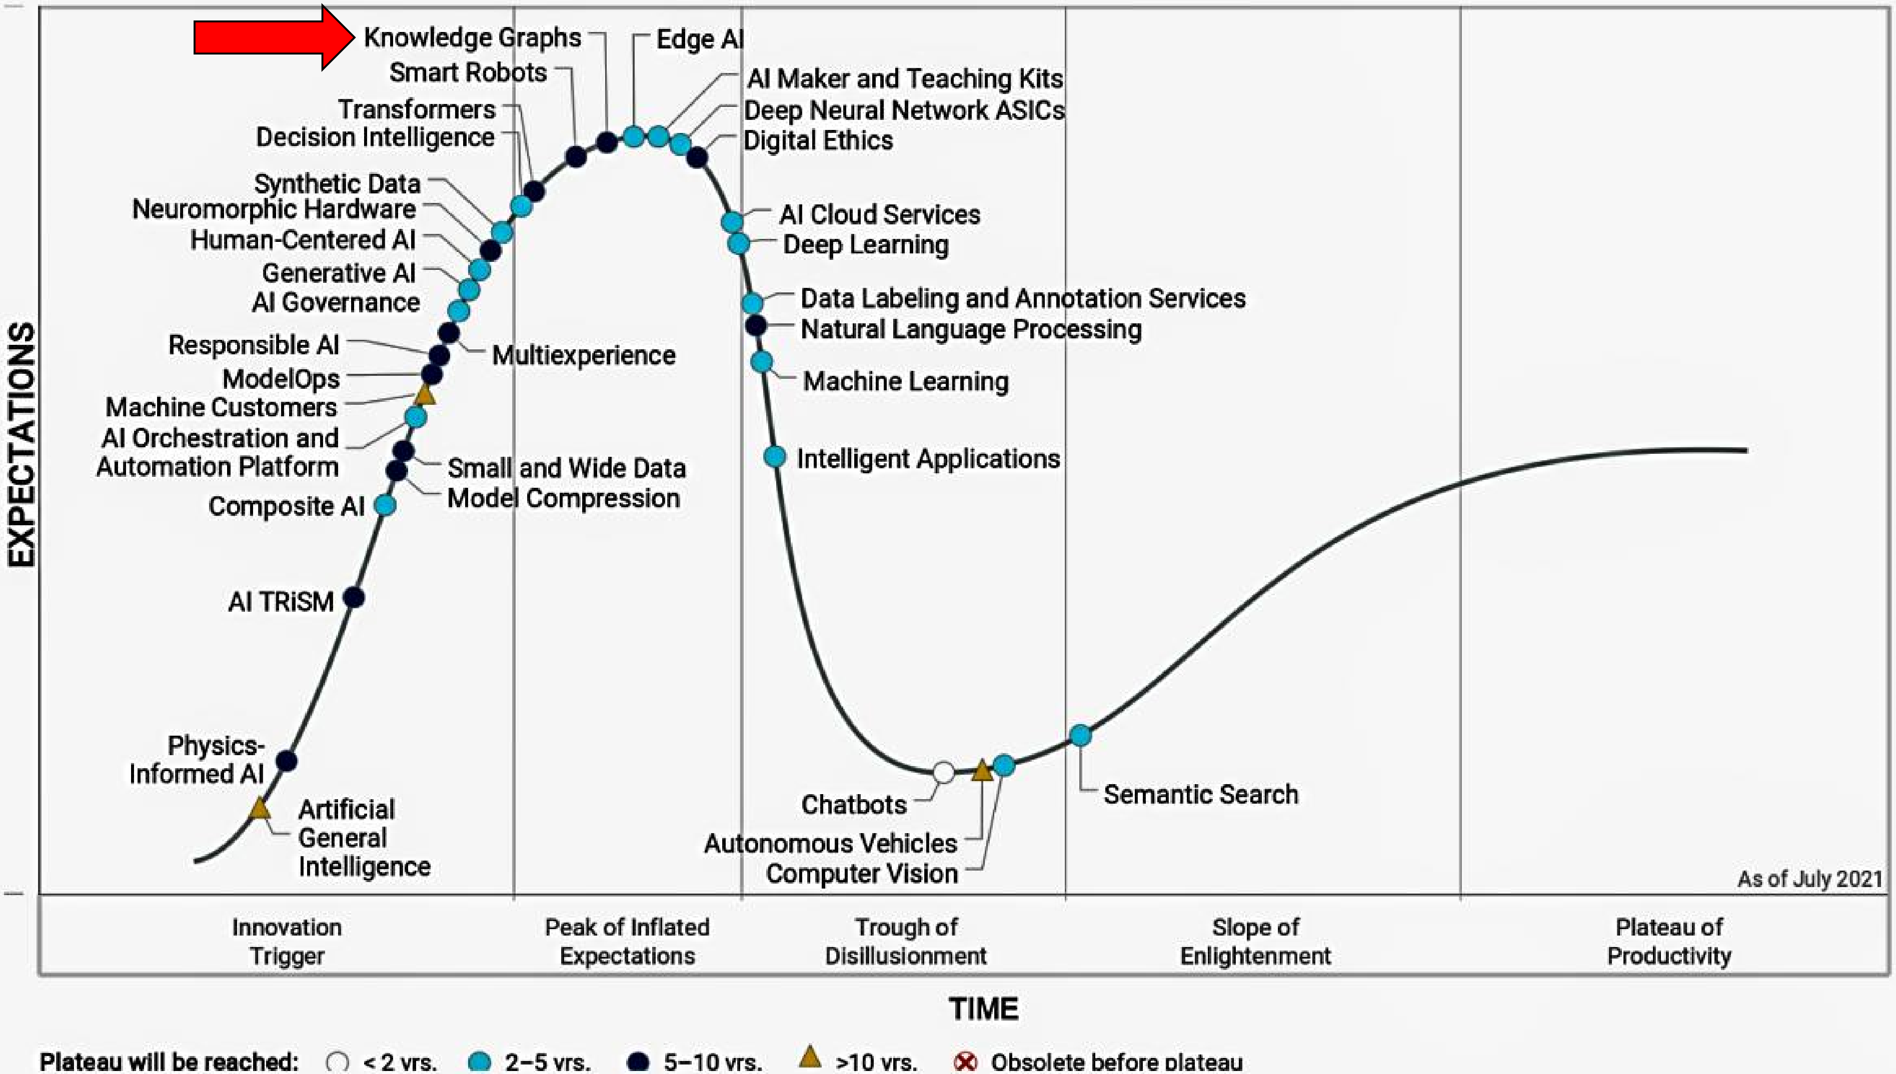

Gartner

Hype Cycle for Artificial Intelligence, 2022

gartner.com

Source: Gartner

Gartner

KANSAS STATE UNIVERSITY

KnowWhereGraph

- **3 years, $5.6M. Follows a $1M, 1-year pilot.**
- **NSF “Open Knowledge Networks” (OKN) program.
21 phase 1 projects; 5 phase 2 projects.**

Team and Partnership

PI: Krzysztof Janowicz, UCSB
**Co-PIs: Mark Schildhauer, Wenwen Li,
Dean Rehberger, Pascal Hitzler**

Public release

- Knowledge Graph with about >12B triples
  - One of the currently largest public knowledge graphs.
  - Focus on spatial data related to environment and natural disasters
- (forthcoming)
  - open source software for access and management

http://knowwheregraph.org/

Thematic Datasets

Dataset Name/ Theme	Source Agency	Key Attributes	Spatial Coverage	Temporal Coverage
Soil Properties	USDA	soil type, farmland class	Targeted regions in US	Current
Wildfires	USGS, USDA, USFS, NIFC	wildfire type, burn severity, num. acres burned, contained date	US	1984–current
Earthquakes	USGS	magnitude, length, width, geometry	Global (mag. over 4.5)	2011-01-01 to 2022-01-18
Climate Hazards	NOAA	injuries, deaths, property damages	US	1950–2022
Expert - Covid-19 Mobility	Direct Relief (DR)	name, affiliation, expertise	Global	2021
Expert - General	KWG, UC System, DR, Semantic Scholar	name, affiliation, expertise with spatiotemporal scopes	Global	unlimited
Cropland Types	USDA	crop types (raster data)	US	2008-2021
Air Qual. Obs.	U.S. EPA	AQI value, CO concentration	US	1980–2022
Smoke Plumes	NOAA	daily smoke plumes extent	US	2010-2022
Climate Observations	NOAA	temperature, precipitation, PDSI, PHSI	US	1950 - 2022
Disaster Declaration	FEMA	designated area, program, amount approved, program designated date	US	1953 - 2022
Smoke Plume Extents	NOAA	Smoke extent	US	2017 - 2022
BlueSky Forecasts	Bluesky	PM10, PM5	US	2022-03-07
Transportation (highway network)	DOT	road type, road length, road sign	US	2014
Public Health	CDC, US Census	below poverty level percent, diabetes age adjusted 20 plus percent, obesity age adjusted 20 plus percent	US	2017
Social Vulnerability	CDC/ATSDR	social vulnerability index	US	2018
Hurricane Tracks	NOAA	max wind speed, min pressure	US	1851-2020

Place-Centric Datasets

Place-Centric Dataset	Defining Authority	Spatial Coverage
S2 Cells	Google	Lvl 9 (Global), Lvl 13 (US),
Global Administrative Regions	University of Berkeley, Museum of Vertebrate Zoology and the International Rice Research Institute	Global
US Federal Judicial District	DoJ, ESRI	US
National Weather Zones	NOAA	US
FIPS Codes	NRCS	US
Designated Market Area	Nielen	US
ZIP	ZCTA	US
Climate Division	NOAA	US
Census Metropolitan Area	US Census	US
Drought Zone	NDMC, USDA,NOAA	US
Geographic Name Information System	USGS	US

DaSe Lab

KANS
UNIV

Contents

- Two current trends:
  - Neuro-Symbolic Artificial Intelligence
  - Knowledge Graphs
- And their convergence:
  - Added Value for Deep Learning
    - Example: Explainable AI
  - Added Value for Knowledge Graphs
    - Example: Deep Deductive Reasoning

Added Value for Deep Learning

Prospects

- KGs are a rich source of structured training data
- KGs are a rich source of background knowledge

- Improved performance and trainability of DL systems
- Interpreting and explaining DL systems via background knowledge

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.

Concept

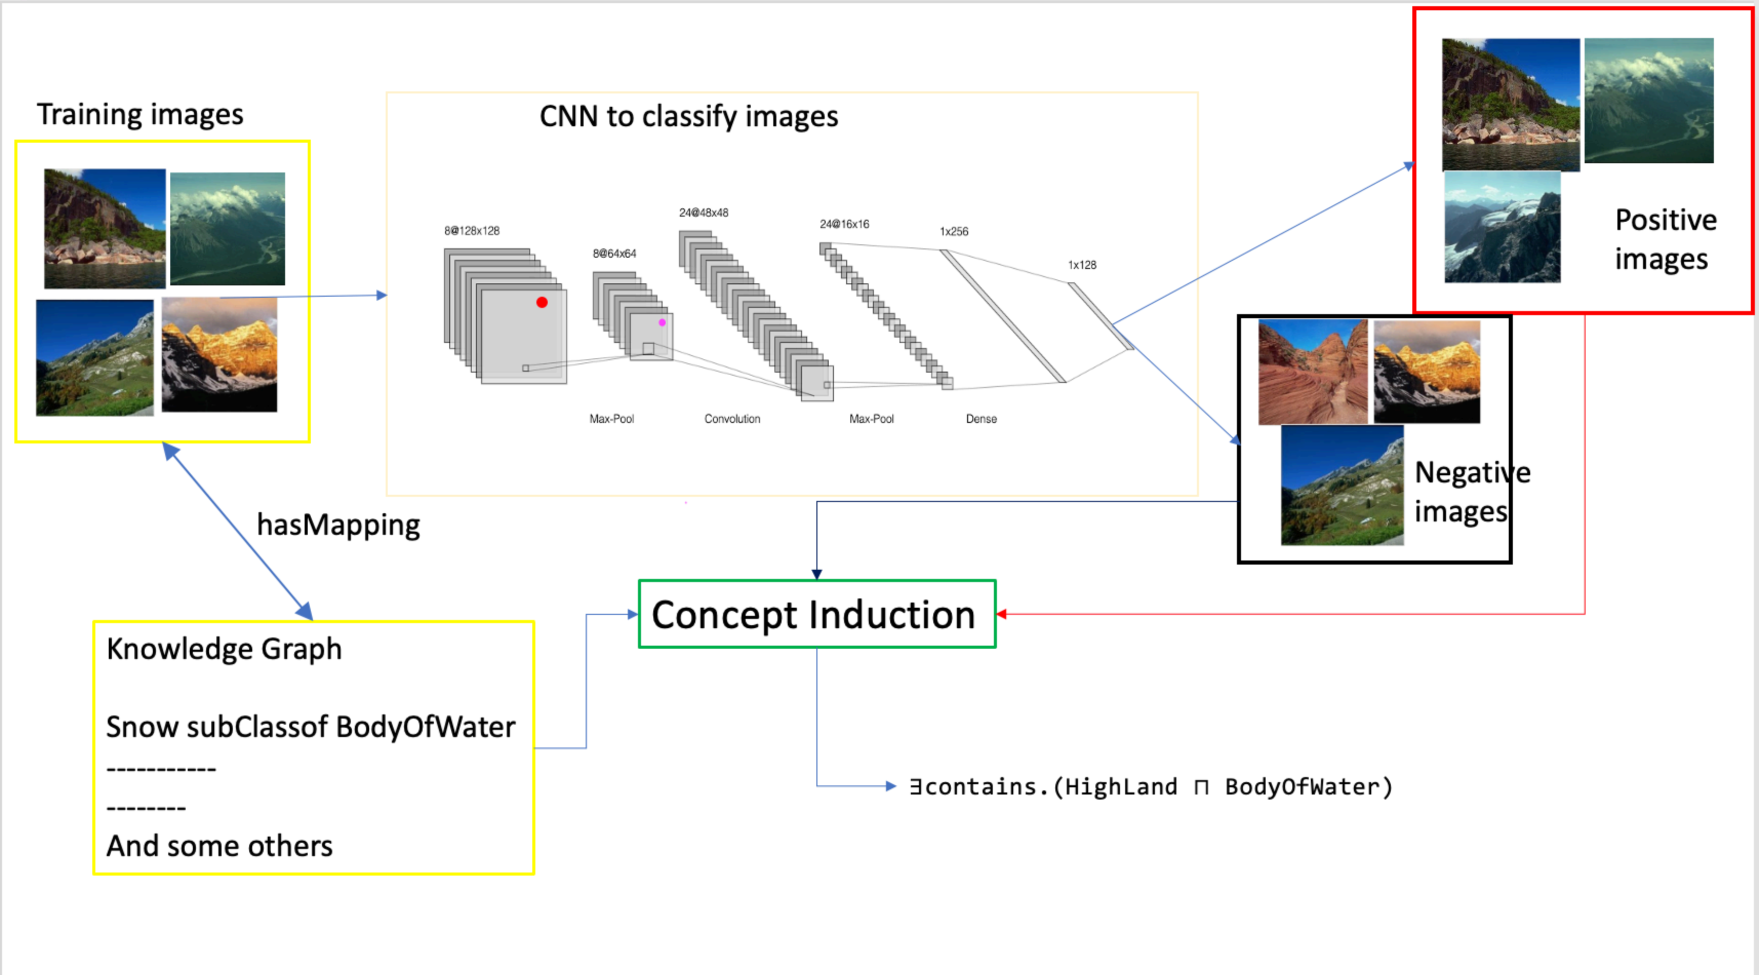

Training images
CNN to classify images
8@128x128
8@64x64
24@48x48
24@16x16
1x256
1x128
Max-Pool
Convolution
Max-Pool
Dense
Positive images
Negative images
hasMapping
Concept Induction
Knowledge Graph
Snow subClassof BodyOfWater

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:

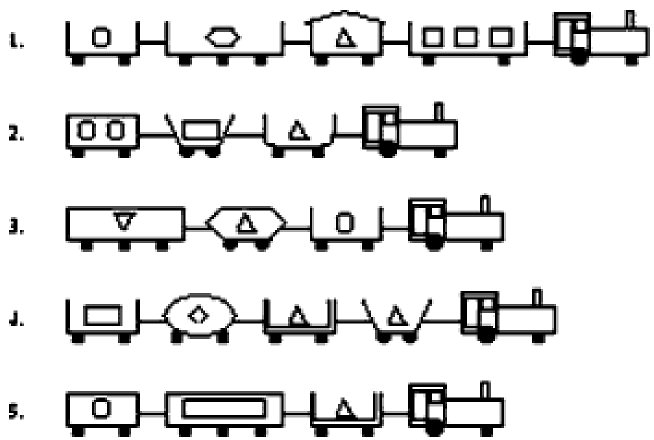

negative examples:

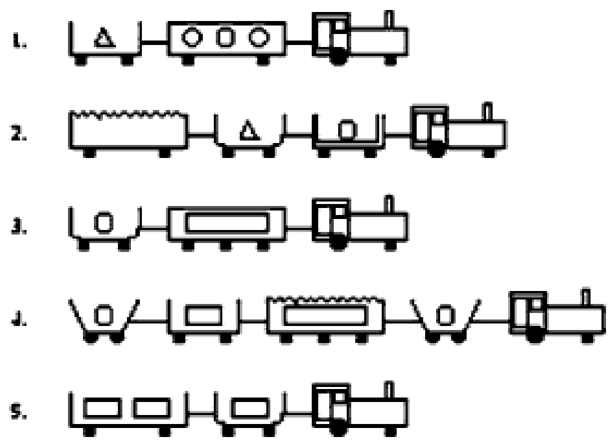

Task: find a class description (logical formula) which separates positive and negative examples.

DL-Learner

Positive examples:

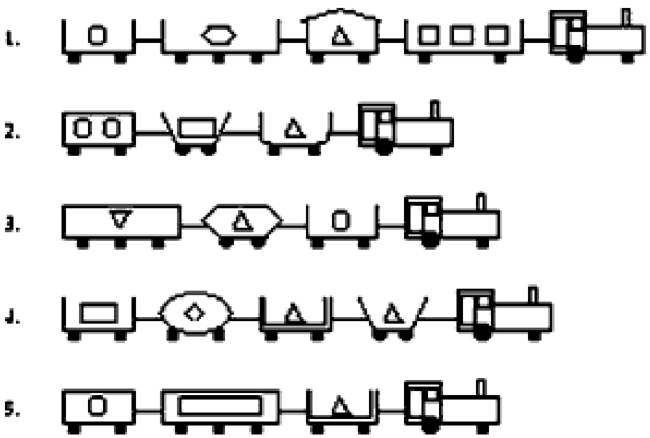

negative examples:

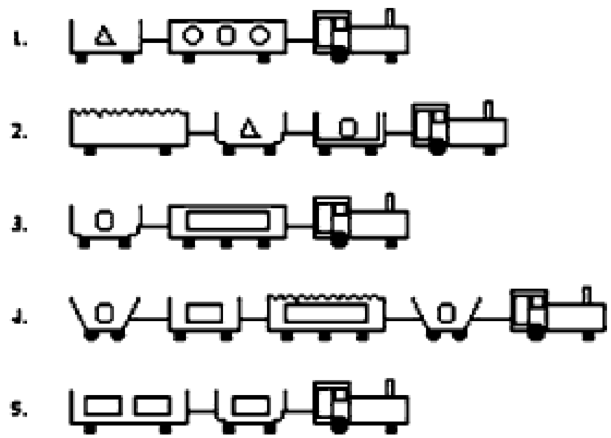

DL-Learner result: $\exists \texttt{hasCar}.(\texttt{Closed} \sqcap \texttt{Short})$

In FOL:

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

ECII: heuristic Concept Induction system

- **For scalability, we developed ECII (Efficient Concept Induction from Instances) which trades some correctness for speed. [Sarker, Hitzler, AAAI-19]**

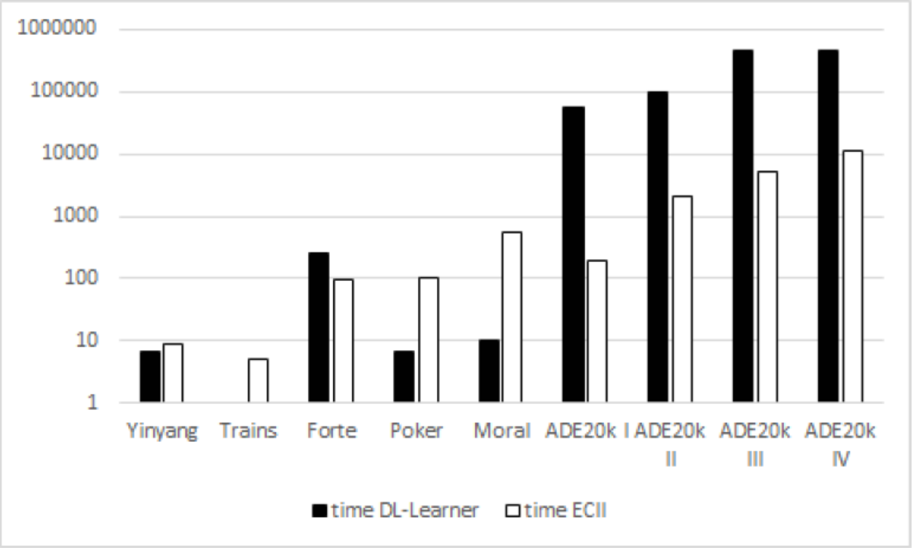

Figure 1: Runtime comparison between DL-Learner and ECII. The vertical scale is logarithmic in hundredths of seconds, and note that DL-Learner runtime has been capped at 4,500 seconds for ADE20k III and IV. For ADE20k I it was capped at each run at 600 seconds.

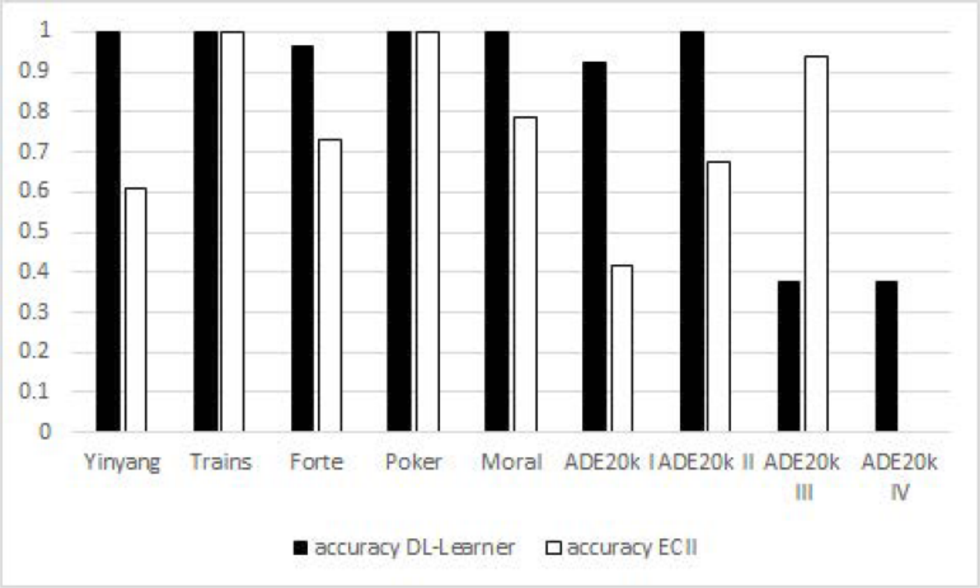

Figure 2: Accuracy (α_3) comparison between DL-Learner and ECII. For ADE20k IV it was not possible to compute an accuracy score within 3 hours for ECII as the input ontology was too large.

Proof of Concept Experiment

DaSe Lab

Positive:

Negative:

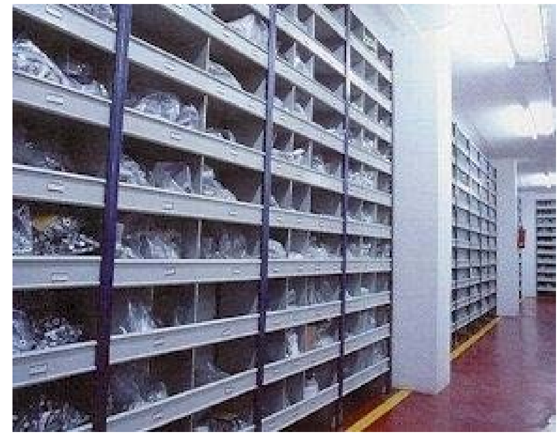

Images

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 Background Knowledge

- Wikipedia category hierarchy (curated)
- approx. 2M concepts
- For each known object in image, create an individual for the ontology which is in the appropriate class.

contains road1

contains window1

contains door1

contains wheel1

contains sidewalk1

contains truck1

contains box1

contains building1

Idea Recap

- Generate explanation of the whole model
- Global explanation

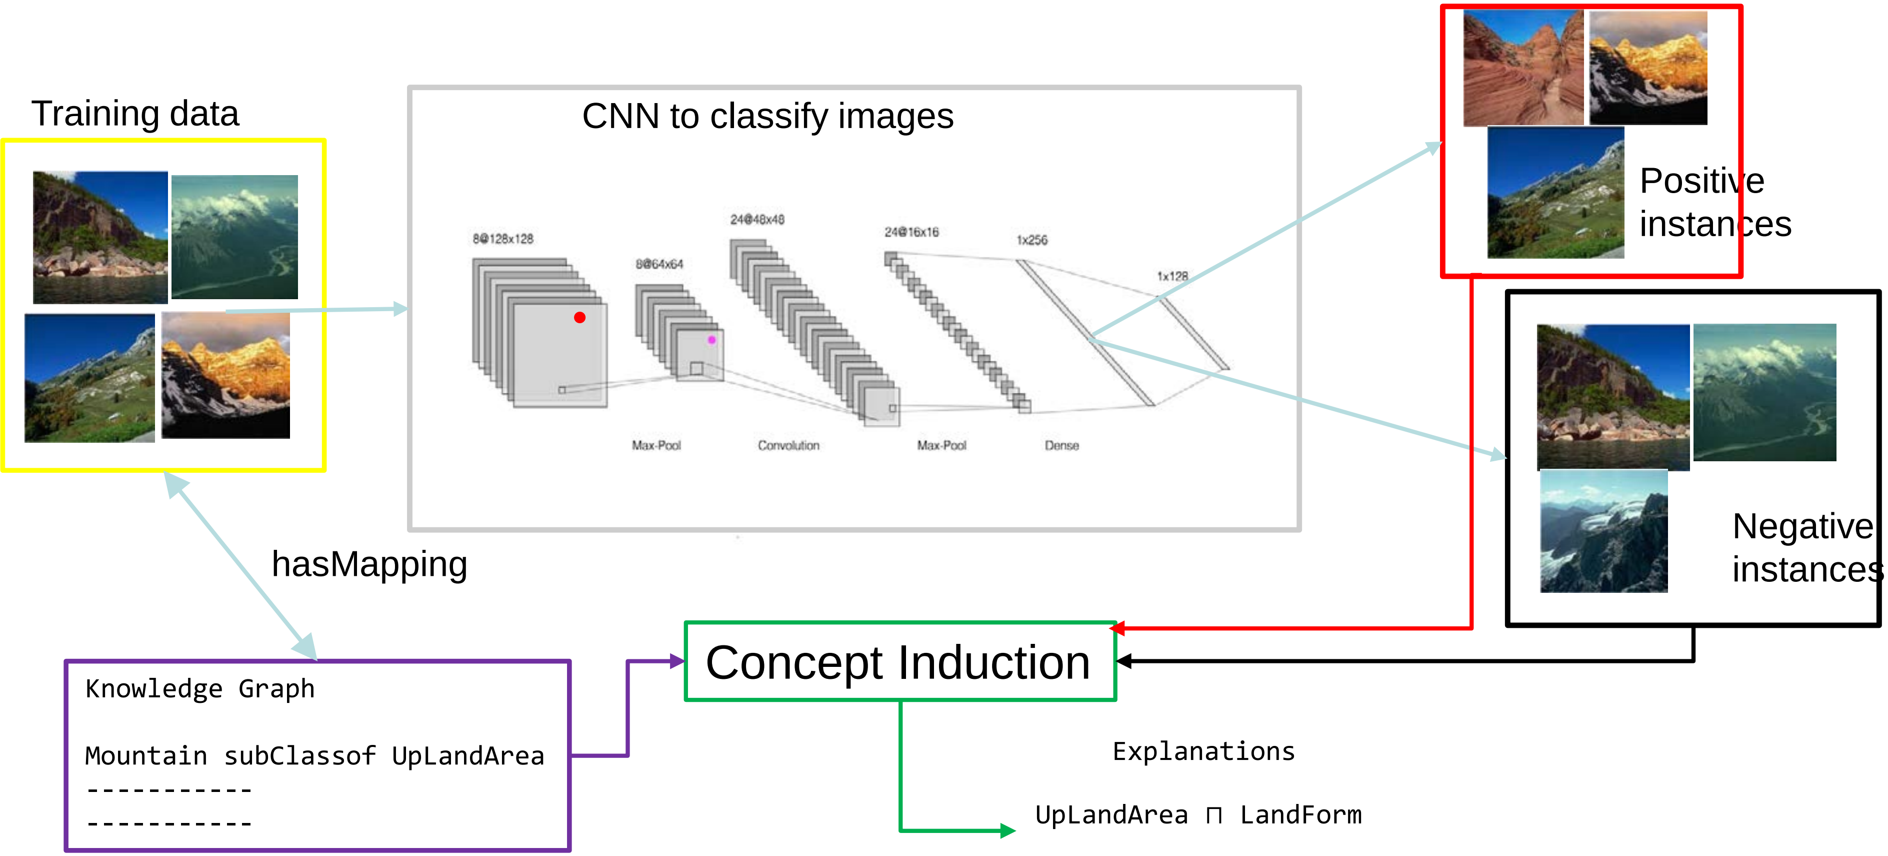

Understanding hidden layer activations Through Concept Induction

Results

Neuron number 04 (dense layer, i.e. before output layer):

- Total number of images that got activated = **612/1370** (1370= test_dataset)
- Highest activation = **12.627778**
- Total number of positives = **149 (images that has value >= 6)**
- Total number of negatives = **150 (images that has value < 6)**

Solution given by ECII analysis for neuron 04

solution 1: (:Bed)
solution 2: (:WN_Bed)
solution 3: (:WN_Table)
solution 4: (:WN_Lamp)
solution 5: ((:WN_Table) ⊓ (:Bed))
solution 6: (:Night_table)
solution 7: (:Cushion)
solution 8: ((:Cushion) ⊓ (:WN_Cushion))
solution 9: (:WN_Shade)
solution 10: ((:Pillow) ⊓ (:WN_Bed))
solution 14: (:WN_Pillow)
solution 17: ((:WN_Cushion) ⊓ (:WN_Lamp))
solution 19: (:WN_Headboard)
solution 24: ((:WN_Lamp) ⊓ (:Pillow))
solution 25: (:WN_Table)

Distinct Concepts from the solution

Bed
Table
Night Table
Lamp
Pillow
Cushion
Headboard

Results

Google analysis for Neuron number 04 :

- Take each concept from distinct concept list for eg: Bed, Table and collect images from Google.
- First set analysis, all images activate (853 images)
- Second set analysis, all images activate (900 images)

ADE20K Dataset

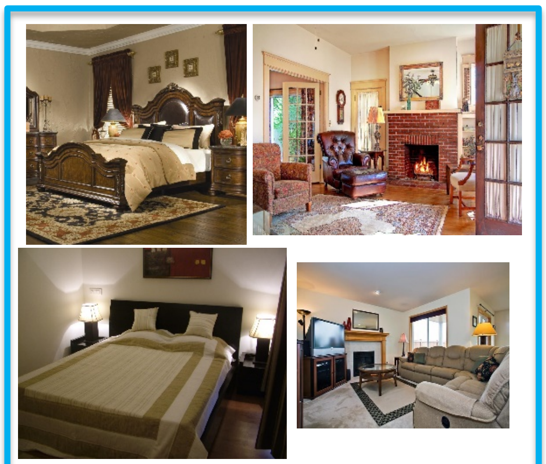

Positive Images

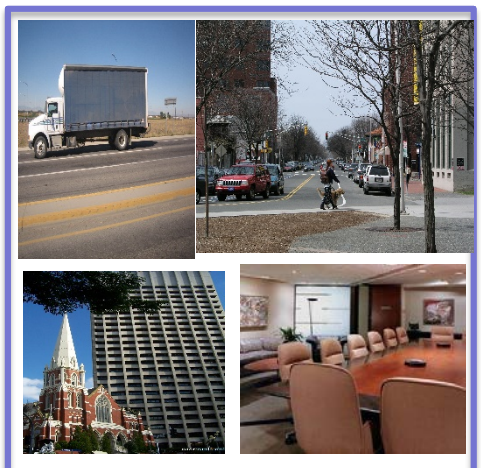

Negative Images

Google Images

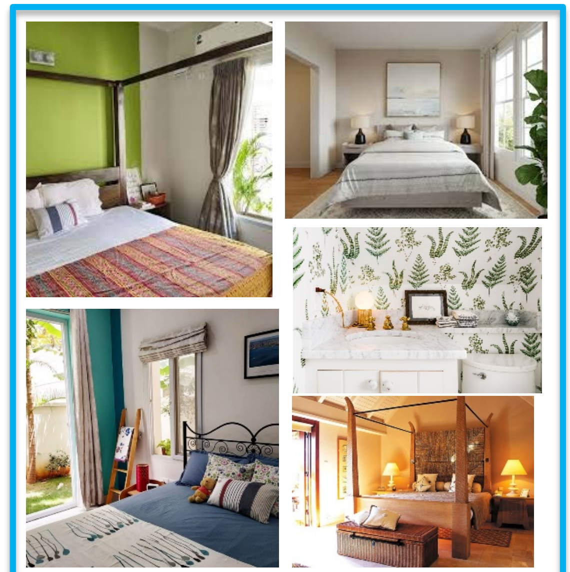

Results

Neuron number 05 :

- Total number of images that got activated = **787/1370** (1370= test_dataset)
- Highest activation = **10.196102**
- Total number of positives = **116 (images that has value >= 5)**
- Total number of negatives = **150 (images that has value < 5)**

Solution given by ECII analysis for neuron 04

solution 1: (:WN_Table)
solution 2: (:Floor)
solution 4: (:WN_Flooring)
solution 5: (:Window)
solution 7: ((:WN_Flooring) ⊓ (:Window))
solution 10: ((:Ceiling) ⊓ (:WN_Table))
solution 15: (:Picture)
solution 17: (:WN_Picture)
solution 22: (:Chair)
solution 24: (:WN_Lamp)
solution 26: ((:WN_Windowpane) ⊓ (:WN_Painting))

Distinct Concepts from the solution

Table
Floor
Window
Ceiling
Picture
Chair
Lamp
Painting

Results

Google analysis for Neuron number 05 :

- Take each concept from distinct concept list for eg: Window, Chair, Picture and collect images from google.
- First set analysis, all images activate (1500 images)
- Second set analysis, all images activate (508 images)

ADE20K Dataset

Positive Images

Negative Images

Google Images

Results

Neuron number 11 :

- Total number of images that got activated = **794/1370** (1370= test_dataset)
- Highest activation = **17.6951**
- Total number of positives = **262 (images that has value >= 9)**
- Total number of negatives = **250 (images that has value < 9)**

Solution given by ECII analysis for neuron 11

solution 1: (:WN_Edifice)
solution 2: (:WN_Building)
solution 3: (:Building)
solution 4: (:WN_Sky)
solution 5: (:Sky)
solution 6: (:WN_Road)
solution 7: (:WN_Route)
solution 8: (:Road)
solution 9: (:WN_Tree)
solution 10: ((:WN_Motorcar) ⊓ (:WN_Machine))
solution 14: (:WN_Automobile)
solution 17: ((:WN_Route) ⊓ (:WN_Building))
solution 19: ((:WN_Automobile) ⊓ (:WN_Route))
solution 24: (:Sidewalk)
solution 25: (:WN_Pavement)

Distinct Concepts from the solution

Edifice(Building)
Building
Sky
Road
Route
Tree
Motorcar
Machine
Automobile
Sidewalk
Pavement

Results

Google analysis for Neuron number 11 :

- Take each concept from distinct concept list for eg: Building, Sky and collect images from google.
- First set analysis, all images activate (183 images)
- Second set analysis, all images activate (454 images)

ADE20K Dataset

Positive Images

Negative Images

Google Images

Are Concept Induction explanations meaningful to humans?

Are the results human-compatible? Part I

- **Hypothesis:**
  - **ECII explanations are better than semi-random explanations, but worse than human-generated explanations.**
- **Experimental setting as before.**
- **300 Amazon Mechanical Turk participants**
- **Seven concepts taken from top ECII results.**
- **45 image set pairs, each set corresponding to a category.**

Are the results human-compatible? Part I

Which of these better represents what the images in group A have that the images in group B do not?

Bake, Bakery, Bread, Indoor, Product, Store, Woman	Basket, Bread, Cake, Ceiling, Floor, Person, Wall

Are the results human-compatible? Part I

DaSe Lab

87-13 **97-3** **87-12**

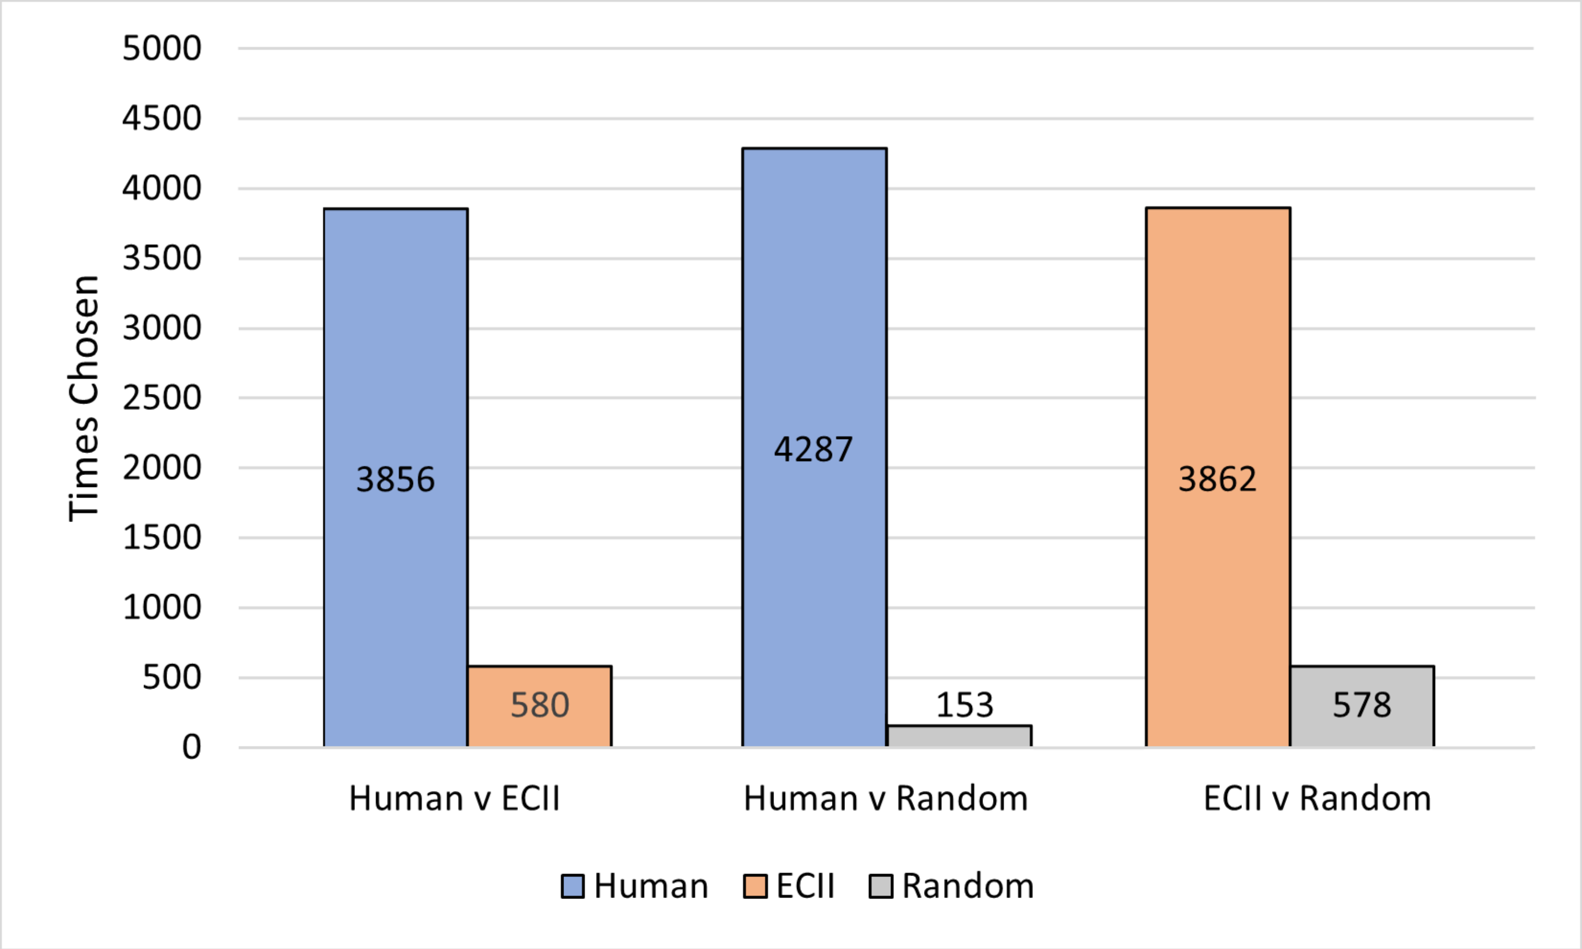

Are the results human-compatible? Part II

- **Hypothesis:**
  - **ECII explanations matched to correct images better than chance, but not as frequently as human generated explanations**
- **Experimental setting as before.**
- **100 Amazon Mechanical Turk participants**
- **16 image sets, from ML decision errors (logistic regression classifier)**

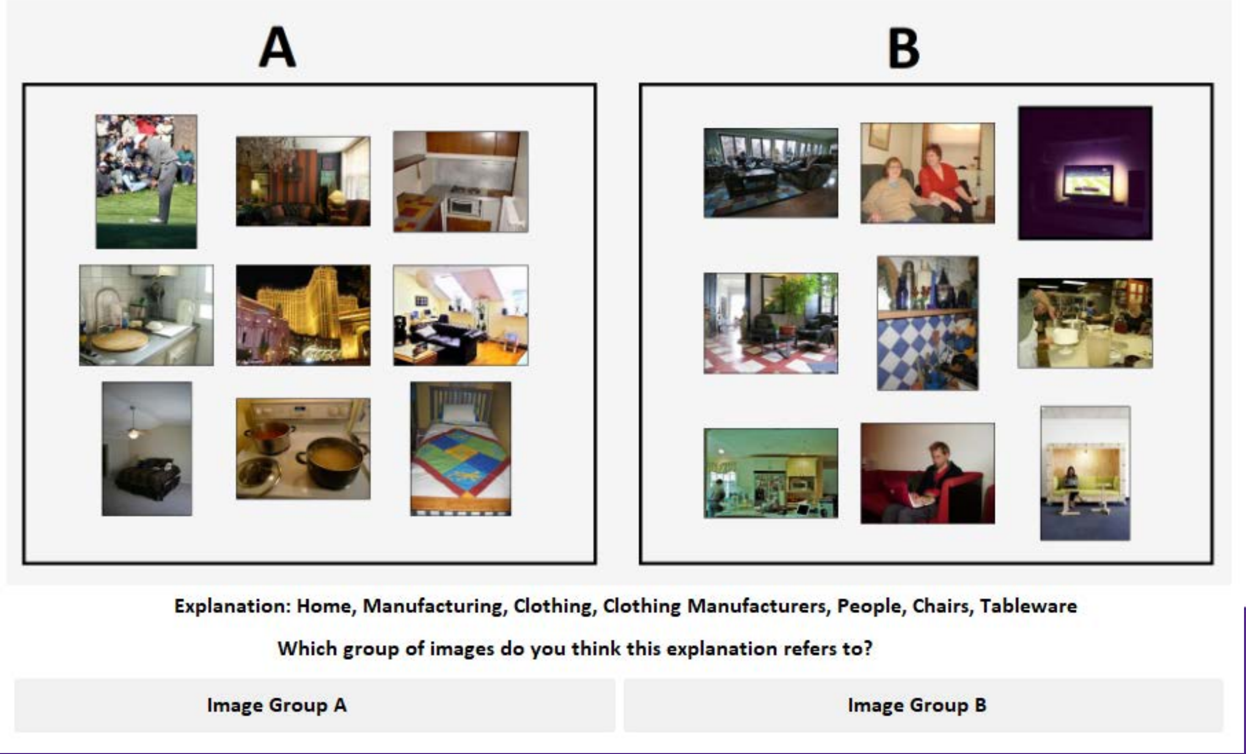

Are the results human-compatible? Part II

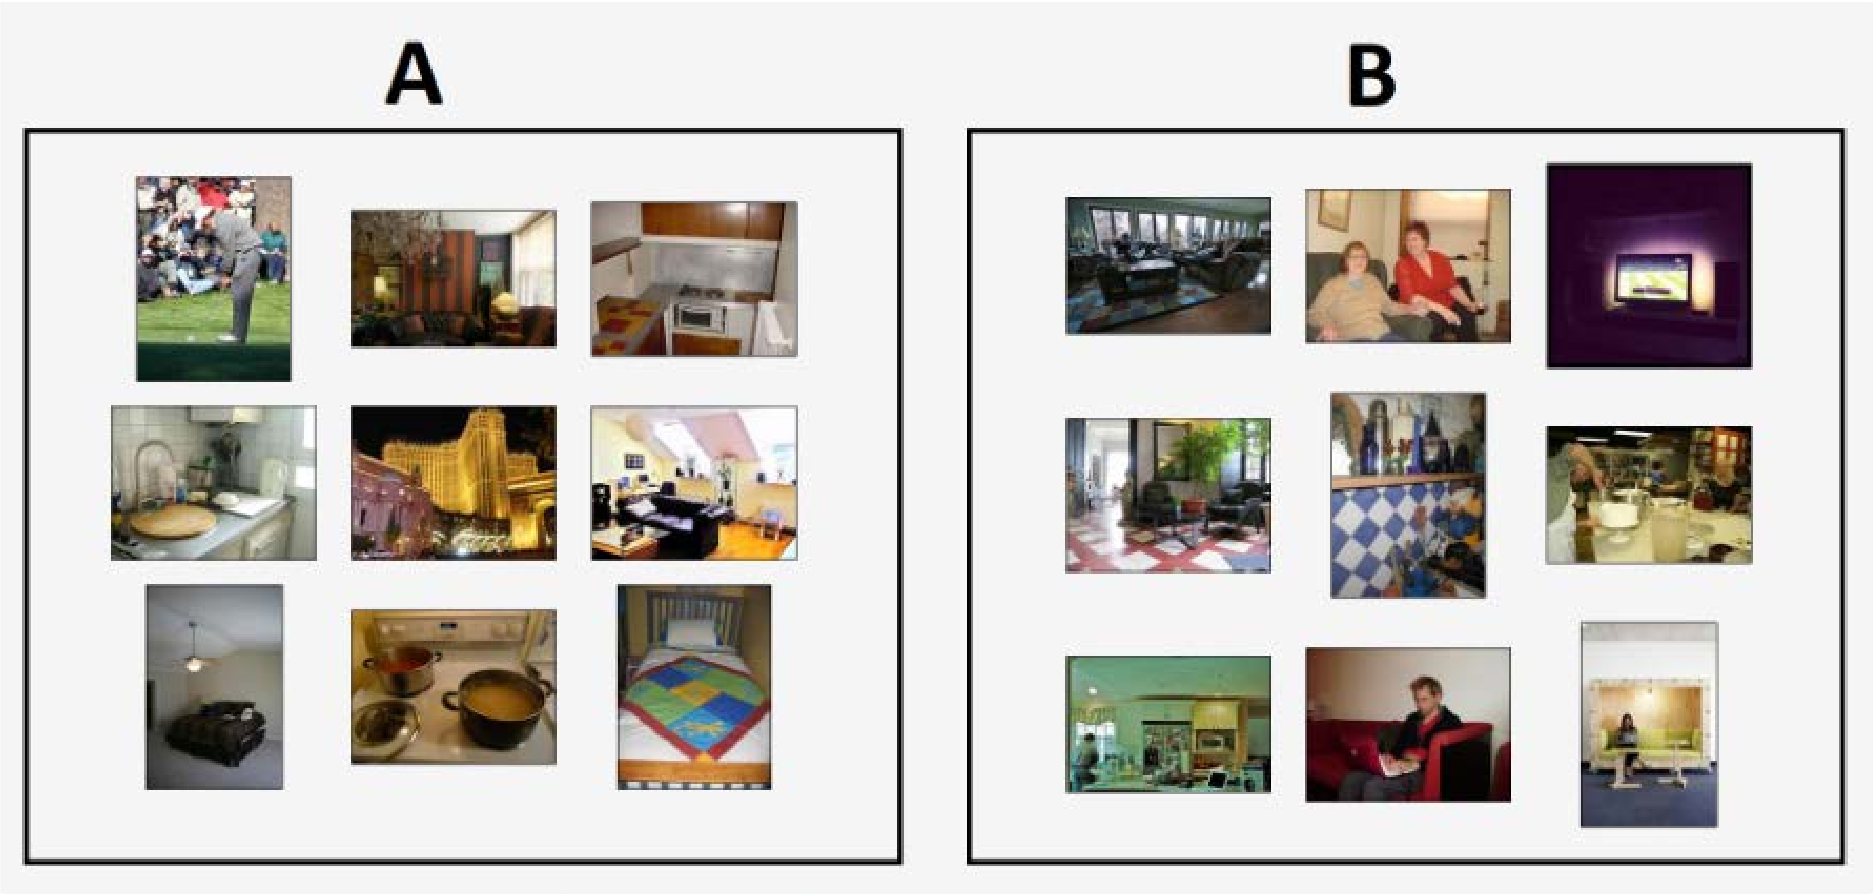

Explanation: Home, Manufacturing, Clothing, Clothing Manufacturers, People, Chairs, Tableware

Which group of images do you think this explanation refers to?

Image Group A	Image Group B

Are the results human-compatible? Part II

- **Bayesian hierarchical signal-detection model (SDT)**
  - **yields discriminability measure**

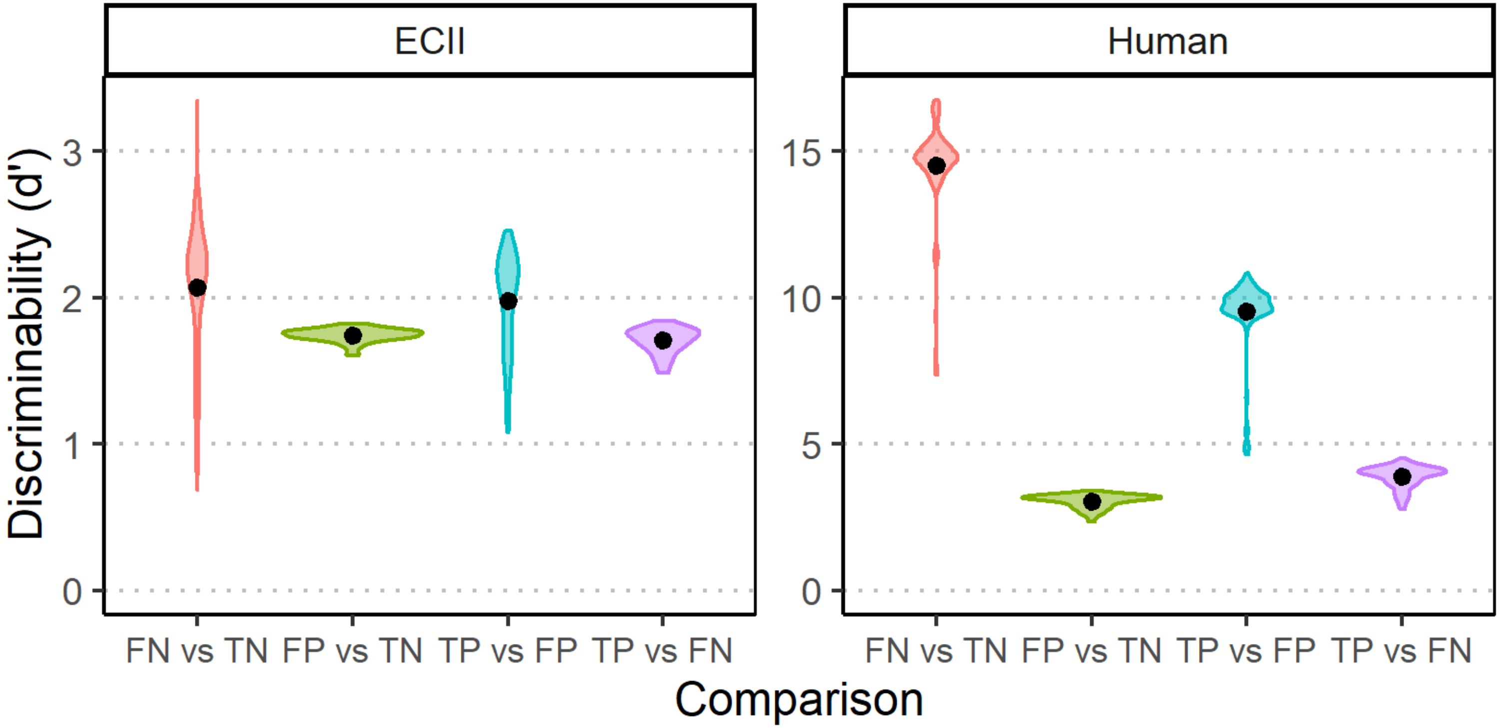

Improving Deep Learning through Concept Induction

Improving deep learning

Experimental set-up

- **Dataset : Twitter Dataset for toxicity analysis**
  - https://www.kaggle.com/competitions/jigsaw-unintended-bias-in-toxicity-classification/data
  - **Classes like "Lie, Dangerous, Insult"**

- **Language Model Used: Bert Base Model**
  - **12 layers**
  - **768 hidden layer neurons**
  - **110M parameters**

Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. NAACL-HLT (1) 2019: 4171-4186

Data examples – “Insult” class

- "Fiore, an occupation sympathizer..." This article makes me feel sick. An insult to Oregonians who have tolerated 41 days and more from this unwanted intrusion. An insult to the LE that put their lives and reputations at risk to resolve this. The mutual admiration between her and Bundy's counsel is to be expected.
  correctly classified
- I'm not sure what you're trying to say, or what the source is of you're information you've implied is somehow not relevant to this article. Forget about mainstream media and the tired and over used commentary that dismiss all mainstream media and politicians making up canned rhetoric repeating it so often that easily manipulated people actually believe them.   We all need to worry about individuals that have an ax to grind and make statements out of thin air, try to shock and change the subject on issues. There is racism in our country and it has been passed down from one generation to another but all good people with moral compasses will continue to work within the process by joining together for the rights of all human beings, we will all benefit and it has nothing to do with political sides blather or insults directed at media. We have options, as a society, our sources for information from credible research is unlimited. You may be looking for truth in all the wrong places.
  incorrectly classified

Concept Induction Analysis

- Run ECII on false positives vs. true positives
- Take first 20 results from ECII
- Get new examples that fall under all of the ECII classes
- Retrain with the additional examples
  - initial training set size: 10,000
  - retraining set size: 11,800
  - i.e. 18% added

Does retraining improve classification?

Results before and after training

Class	Accuracy (before)	**Accuracy (after)**	Precision (before)	**Precision (after)**	F-Measure (before)	**F-Measure (after)**	Recall (before)	**Recall (after)**
Lie	0.9483	**0.9721**	0.9333	**0.9464**	0.9589	**0.9789**	0.9859	**0.9897**
Dangerous	0.8731	**0.8947**	0.8485	**0.8711**	0.8682	**0.8890**	0.8889	**0.9120**
Crazy	0.8911	**0.9105**	0.8511	**0.8784**	0.8791	**0.8962**	0.9090	**0.9465**
Corruption	0.9455	**0.9788**	0.9167	**0.9533**	0.8800	**0.9125**	0.8462	**0.8782**
Fool	0.8983	**0.9427**	0.9483	**0.9788**	0.9016	**0.9433**	0.8594	**0.9652**
Insult	0.7813	**0.8333**	0.7885	**0.8123**	0.7961	**0.8211**	0.8039	**0.8349**

Contents

- Two current trends:
  - Neuro-Symbolic Artificial Intelligence
  - Knowledge Graphs
- And their convergence:
  - Added Value for Deep Learning
    - Example: Explainable AI
  - Added Value for Knowledge Graphs
    - Example: Deep Deductive Reasoning

Added Value for Knowledge Graphs

Prospects

DL systems to assist with

- schema (ontology) modeling
- KG construction based on schema
- schema alignment
- co-reference resolution
- data quality assurance
- KG reasoning

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.**

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

Memory Network based on MemN2N

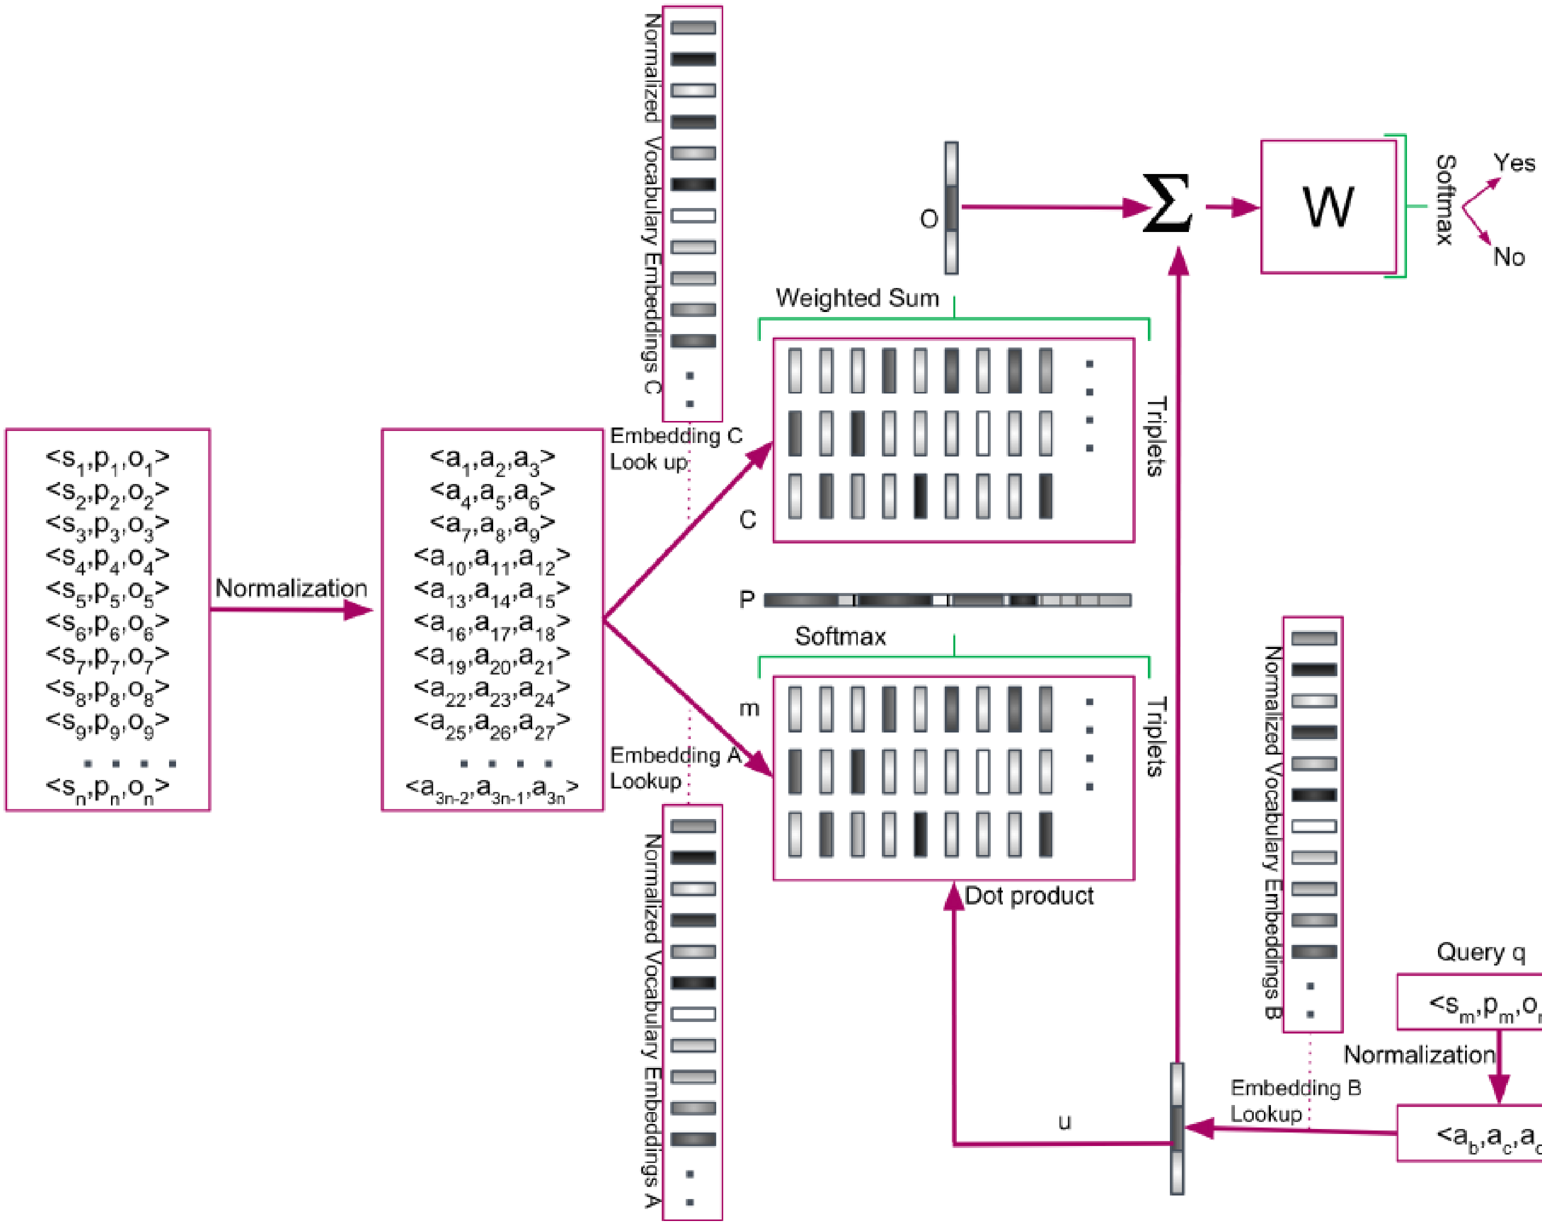

Experiments: Performance

Training	Test	Valid Triples Class			Invalid Triples Class			Accuracy
		Prec (%)	Rec	F-Measure	Prec	Rec	F-Measure	
A	LD 1	93	98	96	98	93	95	**96**
A (90%)	A (10%)	88	91	89	90	88	89	**90**
A	B	79	62	68	70	84	76	**69**
A	Synth 1	65	49	40	52	54	42	**52**
A	LD 2	54	98	70	91	16	27	86
C	LD 2	62	72	67	67	56	61	91
C (90%)	C (10%)	79	72	75	74	81	77	80
A	D	58	68	62	62	50	54	58
C	D	77	57	65	66	82	73	73
A	Synth 2	70	51	40	47	52	38	51
C	Synth 2	67	23	25	52	80	62	50

Baseline: non-normalized embeddings, same architecture

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

[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

Conclusions

Conclusions

- Two current trends:
  - Knowledge Graphs
  - Neurosymbolic AI
- Plenty of opportunities
  - Improving DL systems with KG-based background knowledge
    - Explainable AI by Concept Induction
  - Solving key KG problems using DL approaches.
    - Deep Deductive Knowledge Graph Reasoning

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 34 (3), 197-209, 2022.

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)

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.

Cara Widmer, Md Kamruzzaman Sarker, Srikanth Nadella, Joshua Fiechter, Ion Juvina, Brandon Minnery, Pascal Hitzler, Joshua Schwartz, Michael Raymer, Towards Human-Compatible XAI: Explaining Data Differentials with Concept Induction over Background Knowledge. arXiv:2209.13710

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.

Henry Kautz, The third AI summer: AAAI Robert S. Engelmore Memorial Lecture, AI Magazine 43, 2022, 105-125

Thanks!