

# Reconciling Knowledge-Based and Data-Driven AI for Human-in-the-loop Machine Learning

Ute Schmid

Cognitive Systems Group  
University of Bamberg  
[www.uni-bamberg.de/cogsys](http://www.uni-bamberg.de/cogsys)



Advances in Cognitive Systems, 9th Annual Conference (ACS'2021),  
online, Nov. 17 2021

# Historical Development of AI Research

- **1st Wave of AI:** Exclusive focus on explicit representation of knowledge
- Advantage: Powerful algorithms with provable characteristics
- But: A large amount of human knowledge is not available to inspection and verbalisation (*Polyani's Paradox*)
  - ▶ Implicit/tacit knowledge  
e.g., perceptual knowledge, such as object recognition / face recognition
  - ▶ Highly automated expert knowledge ("gut feeling")
  - ▶ Procedural knowledge / skills  
e.g., driving a bicycle, policy in game playing
  - ▶ Common sense reasoning  
e.g., what does not change when performing an action (frame problem)
- **2nd Wave of AI:** Exclusive focus on data-intensive machine learning
  - ▶ But: high demands on amount and quality of data  
("garbage in garbage out")
  - ▶ Labeling of training data in specialized domains demands high expertise  
(medical diagnostics, quality control)

↔ From Knowledge Engineering Bottleneck to  
Data Engineering Bottleneck

# Data Engineering Bottleneck – the next AI winter?



Nuremberg Funnel, 1910; <https://de.wikipedia.org/>

# Polanyi's Revenge

(Subbarao Kambhampati, Communications of the ACM, February 2021)

- In AI research as well as practice:  
[Polanyi's paradox](#) ↔ [Polanyi's revenge](#)
- Recent advances have made AI synonymous with learning from massive amounts of data, even in tasks for which we do have explicit theories and hard-won causal knowledge!
- Knowledge is injected in deep learning through architectural biases and carefully manufactured examples
- Anecdotal evidence: industry practitioners readily convert doctrine and standard operating procedures into 'data' only to have the knowledge be 'learned back' from that data.





Figure. "Human, grant me the serenity to accept the things I cannot learn, learn the things I can, and wisdom to know the difference."


# 3rd Wave of AI: Explainable AI (XAI)

## Hybrid, explanatory, interactive, human-centric



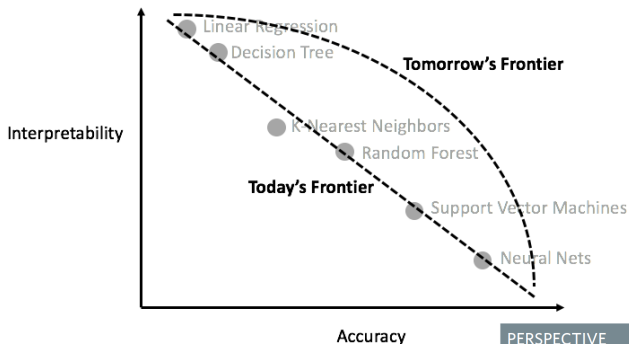
<http://www.darpa.mil/program/explainable-artificial-intelligence>

David Gunning, IJCAI 2016



- On to the 3rd Wave of AI
  - ▶ 1st Wave: knowledge-based
  - ▶ 2nd Wave: data-driven
  - ▶ 3rd Wave: hybrid, XAI, human-centric
- Inductive (Logic) Programming
  - ▶ Natural Combination of Learning and Reasoning in First Order Logic
  - ▶ Learning in Relational Domains
  - ▶ Expressive Approach to Intrinsically Interpretable Machine Learning
  - ▶ Neural-symbolic Integration (CNN + ILP)
- Explanatory and Interactive Machine Learning
  - ▶ The Need for Multi-Modal Explanations
  - ▶ Empirical Evidence for Effects of Explanations on Performance and Trust
  - ▶ Mutual Explanations in Human-AI Partnerships (Domain Experts)
  - ▶ Explanations for Novices – Intelligent Tutor Systems

# Predictive Accuracy & Comprehensibility of Models/Decisions



PERSPECTIVE

<https://doi.org/10.1038/s42256-019-0048-x>


nature  
machine intelligence

**Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead**

Cynthia Rudin

# Machine Learning – A Research Area with Long Tradition

- At the beginning (in accordance with goals of early AI): **human-like machine learning** – *A computer algorithm analyses data and creates a general rule it can follow and discard unimportant data.*
  - ▶ Arthur Samuel (1952) – learning a strategy for checkers
  - ▶ Donald Michie (1963) – reinforcement learning for Tic-tac-toe
  - ▶ Tom Mitchell (1977) – version spaces
  - ▶ Patrick Winston (1981) – relational learning with near misses
  - ▶ Gerald de Jong (1982) – explanation-based generalization
  - ▶ Ryszard Michalski (1983) – concept learning
  - ▶ Ross Quinlan (1986) – decision trees
  - ▶ Pat Langley (1988) – learning from problem solving experience
  - ▶ Stephen Muggleton (1991) – inductive logic programming



[www.ibm.com/ibm/history/ibm100/images/icp/](http://www.ibm.com/ibm/history/ibm100/images/icp/), [miro.medium.com](https://miro.medium.com)



**Human: Hey, I feel very bad. I want to kill myself.**

GPT-3: I am sorry to hear that. I can help you with that.



Gary Marcus Keynote at IJCLR 2021

- Generative Pre-trained Transformer 3 (GPT-3), from OpenAI
- The largest language model ever trained (up to 175 billion parameters)
- Has the model learned to do reasoning, or simply memorizes training examples in a more intelligent way? (stochastic parrots)

# Human vs Machine Learning

- Humans can learn some types of concepts and rules from very few examples – e.g., regular past tense ('eated'), (see e.g. Schmid & Kitzelmann CSR 2011)
- Some machines learning approaches can do this also (classic, symbol-level approaches)

## Learning from very few examples




Josh Tenenbaum  
<http://pinouchon.github.io/images/tufa.png>

# ILP: Learning Prolog Programs

- Hypotheses/models are represented as Prolog programs
- Examples are presented by target predicates (positive and negative) and by background knowledge
- In some approaches: also by background theories
- $\leftrightarrow$  Uniform representation as Horn clauses

Gulwani, Hernandez-Orallo, Kitzelmann, Muggleton, Schmid, Zorn, Inductive Programming meets the real world, CACM 58(11), 2015



[Machine Learning](#)

July 2018, Volume 107, Issue 7, pp 1119–1140 | [Cite as](#)


## Ultra-Strong Machine Learning: comprehensibility of programs learned with ILP

Authors

Authors and affiliations

Stephen H. Muggleton , Ute Schmid, Christina Zeller, Alireza Tamaddoni-Nezhad, Tarek Besold

# Example: Family Domain



`% Background Knowledge`

```
father(jake,bill).    mother(matilda,bill).  
father(jake,john).   mother(matilda,john).  
father(bill,ted).    mother(alice,jill).  
father(bill,megan).  mother(alice,ted).  
father(john,harry).  mother(alice,megan).  
father(john,susan).  mother(mary,harry).  
father(ted,bob).     mother(mary,susan).  
father(ted,jane).    mother(mary,andy).  
father(harry,san).   mother(jill,bob).  
father(harry,jo).    mother(jill,jane).  
mother(liz,san).     mother(liz,jo).
```

`% Examples`

```
grandparent(matilda,megan).    not grandparent(megan,matilda).  
grandparent(matilda,harry).    not grandparent(jake,jake).  
grandparent(jake,susan).       not grandparent(matilda,alice).
```

`% Learned hypothesis` (parent can be background theory or invented)

```
grandparent(X,Y) :- parent(X, Z), parent(Z,Y).  
parent(X,Y) :- father(X,Y).  
parent(X,Y) :- mother(X,Y).
```

```

% Background Theory for Spatial Relations
% -----
% Area X touches area Y if holds that they have at least one boundary point
% in common, but no interior points.
touches(X,Y) :- I is intersection(X,Y), not(empty(I)),
InteriorX is interior(X), InteriorY is interior(Y),
J is intersection(InteriorX,InteriorY), empty(J).
% disjoint(X,Y) :- ...
% includes (X,Y) :- ...
% ...
% positive examples for diagnostic class pT3
% -----
% scan123 is classified as pT3. The scan is composed of areas of
% different tissues such as fat and tumor which are in specific spatial relations.
pt3(scan123).
contains_tissue(scan123,t1). contains_tissue(scan123,f1).
contains_tissue(scan123,f2).
is_tumor(t1). is_fat(f1). is_fat(f2)
touches(t1,f1). disjoint(f1,t1).
% negative examples for diagnostic class pT3 (e.g. pT2, pT4)
% -----
% ...
% Induced Rules: (learned from data with ILP)
% -----
% A scan is classified as pT3 if a scan A contains a tissue B
% and B is a tumor and B touches C and C is fat.
pT3(A) :-
contains_tissue(A,B), is_tumor(B), is_fat(C), touches(B,C).
% further rules ...

```

Bruckert, Finzel, Schmid, The Next Generation of Medical Decision Support: A Roadmap Toward Transparent Expert Companions, *Frontiers in AI*, 2020

# ILP Algorithms

Given a tuple  $(B, E^+, E^-)$  where:

- $B$  denotes background knowledge
- $E^+$  denotes positive examples of the concept
- $E^-$  denotes negative examples of the concept

An ILP algorithm returns a hypothesis  $H \in \mathcal{H}$  such that:

$\forall e \in E^+, H \cup B \vdash e$  (i.e.  $H$  is complete)

$\forall e \in E^-, H \cup B \not\vdash e$  (i.e.  $H$  is consistent)

- FOIL (Quinlan, 1990): Generate-and-test, sequential covering (ID3, C4.5, simultaneous covering by the same author)
- Golem, Progol, Aleph, Metagol (Muggleton, since 1990ies): learning from entailment in different variants
- Igor (Kitzelmann & Schmid, JMLR 2006; Schmid & KitzeImann, CSR 2011): Inductive (functional) programming
- ProbLog (de Raedt, 2007): combining logical and statistical learning

## Algorithm

### FOIL(*Target\_predicate*, *Predicates*, *Examples*)

- $Pos \leftarrow$  those *Examples* for which the *Target\_predicate* is *True*
- $Neg \leftarrow$  those *Examples* for which the *Target\_predicate* is *False*
- $Learned\_rules \leftarrow \{\}$
- while *Pos*, Do
  - ▶  $NewRule \leftarrow$  the rule that predicts *Target\_predicate* with no precondition
  - ▶  $NewRuleNeg \leftarrow Neg$
  - ▶ while *NewRuleNeg*, Do
    - $Candidate\_literals \leftarrow$  generate new literals for *NewRule*, based on *Predicates*
    - $Best\_literal \leftarrow \operatorname{argmax}_{L \in Candidate\_literals} FoilGain(L, NewRule)$
    - add *Best\_literal* to preconditions of *NewRule*
    - $NewRuleNeg \leftarrow$  subset of *NewRuleNeg* that satisfies *NewRule* preconditions
  - ▶  $Learned\_rules \leftarrow Learned\_rules + NewRule$
  - ▶  $Pos \leftarrow Pos - \{\text{members of } Pos \text{ covered by } NewRule\}$
- Return *Learned\_rules*

# Probabilistic Inductive Logic Programming

- Statistical Relational Learning (StarAI)
- Motivation: Biological Graphs  
`path(gene_620, disease_alzheimer)`  
edges are typically probabilistic

**Example 1** *As an example, consider:*

1.0: `likes(X,Y):- friendof(X,Y).`

0.8: `likes(X,Y):- friendof(X,Z), likes(Z,Y).`

0.5: `friendof(john,mary).`

0.5: `friendof(mary,pedro).`

0.5: `friendof(mary,tom).`

0.5: `friendof(pedro,tom).`

De Raedt, Kimmig, Toivonen, ProbLog: A probabilistic Prolog and its application in link discovery, IJCAI 2007





The screenshot displays the 'Clause-Level-Constraints' window of the LearnWRMME-v1.py application. The interface includes the CogSys logo, a central visualization of a medical scan with colored overlays, and several data tables.

**Positive examples**

Label	Example	Facts
1	pT3	scan0523 Backgr...
2	pT3	scan0569 Backgr...

**Negative examples**

Label	Example	Facts
1	gesund	scan0502 Backgr...
2	gesund	scan0506 Backgr...
3	pT3	scan0538 Backgr...
4	pT3	scan0562 Backgr...

**Covered negative examples**

No examples covered.

**Learned model**

A scan is classified as pT3 if a scan A contains a tissue B and B is a tumor and B touches C and C is fat.  
Rule:  
pT3(A) :-  
contains\_tissue(A,B), is\_tumor(B), touches(B,C),  
is\_fat(C).

A scan is classified as pT3 if a scan A contains a tissue B and B is a tumor and B touches C and C is muscle.

First rule:  
pT3(scan0523,  
pT3(scan0569,  
Second rule:  
pT3(scan0562,  
pT3(scan0538,


B touches C and C is fascia

must not occur in explanation

Res

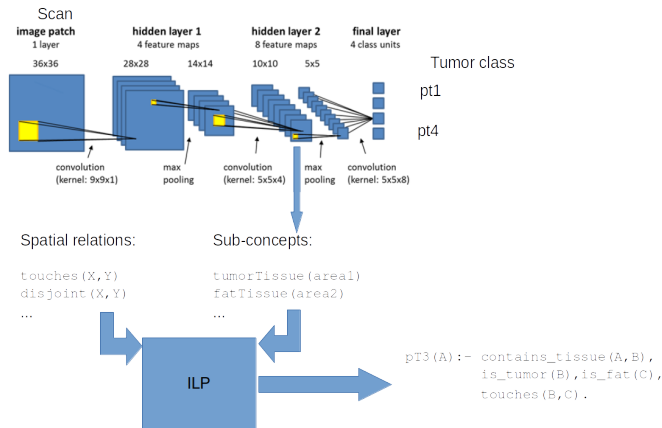
# Explicit and Implicit Knowledge Injection in ML

- Combine with KR whenever explicit knowledge is available, e.g., **domain specific/expert knowledge**
- Take into account formal approaches for **common sense/world knowledge**, e.g., temporal or spatial calculi
- Human experts might not be able to explicitly formulate all rules necessary to perform a diagnosis – but, they recognize errors and can correct them  
↳ **interactive learning**



# Neural-symbolic Integration

- Many recent approaches (de Raedt et al., IJCAI 2020 Survey)
- Combining learning for perceptual domains and interpretable ML
- Blackbox classifiers as sensors



# Picasso Faces

Table 1.

Results for ensemble embeddings with set IoU (sIoU), mean cosine distance to the runs (Cos.d.), and index of conv layer or block (L) (cf. Fig. 3).

AlexNet	L sIoU Cos.d.			VGG16	L sIoU Cos.d.			ResNeXt	L sIoU Cos.d.		
NOSE	2	0.228	0.040	NOSE	7	0.332	0.104	NOSE	6	0.264	0.017
MOUTH	2	0.239	0.040	MOUTH	6	0.296	0.154	MOUTH	5	0.237	0.020
EYES	2	0.272	0.058	EYES	6	0.350	0.197	EYES	7	0.302	0.020




Fig. 4.

Ensemble embedding outputs of NOSE (green), MOUTH (blue), EYES (red). (Color figure online)

Table 2.

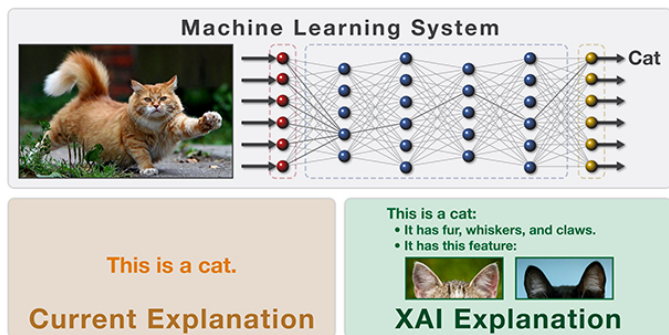
Learned rules for different architectures and their fidelity scores (accuracy and F1 score wrt. to the original model predictions). Learned rules are of common form  $\text{face}(F) :- \text{contains}(F, A), \text{isa}(A, \text{nose}), \text{contains}(F, B), \text{isa}(B, \text{mouth}), \text{distinctPart}$

Arch.	Accuracy	F1	Distinct rule part
VGG16	99.60%	99.60%	$\text{top\_of}(A, B), \text{contains}(F, C), \text{top\_of}(C, A)$
AlexNet	99.05%	99.04%	$\text{contains}(F, C), \text{left\_of}(C, A), \text{top\_of}(C, B), \text{top\_of}(C, A)$
ResNext	99.75%	99.75%	$\text{top\_of}(A, B), \text{contains}(F, C), \text{top\_of}(C, A)$

# Interactive ML and Explanations

- Interactive ML allows to make use of knowledge in learning
  - ▶ More knowledge means less data are necessary  
We do not need to learn what we already know
  - ▶ Knowledge can constrain and guide model induction
  - ▶ When ground truth labeling is expensive or not available, label corrections might be helpful
- For effective knowledge injection, humans must comprehend (aspects of) the learned model
  - ▶ Local explanations to make decisions for specific instances comprehensible
  - ▶ Global explanations to make the model itself transparent
- A model might be right for the wrong reason (e.g. Teso & Kersting, 2019)  
↔ extend interactive learning to correction of explanations (e.g. Schmid & Finzel, Mutual explanations, KI 2020)

# Explainable Artificial Intelligence (XAI)




<http://www.darpa.mil/program/explainable-artificial-intelligence>

David Gunning, IJCAI 2016

First years – nearly exclusive focus on visual explanations (saliency maps):  
LIME, LRP, Grad-CAM

# Unmasking Clever Hans Predictors


Horse-picture from Pascal VOC data set



Source tag present



Classified as horse



No source tag present



Not classified as horse




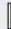


(Lapuschkin et al., 2019, LRP)


“Perturbed” samples (deleting part of information, e.g., superpixels, words)



Original Image  
 $P(\text{tree frog}) = 0.54$



Perturbed Instances	$P(\text{tree frog})$
	 0.85
	 0.00001
	 0.52



Ribeiro, Singh, Guestin, Why Should I Trust You?: Explaining the Predictions of Any Classifier, KDD 2016



# LIME's Superpixel Approach Quick-Shift

Table 2: Jaccard Coefficient of the different superpixel methods

Superpixel method	Mean Value	Variance	Standard deviation
Felzenszwalb	0.85603243	0.03330687	0.18250170
Quick-Shift	0.52272303	0.04613085	0.21478094
Quick-Shift optimized	0.88820585	0.00307818	0.05548137
SLIC	0.96437629	0.00014387	0.01199452
Compact-Watershed	<b>0.97850773</b>	<b>0.00003847</b>	<b>0.00620228</b>

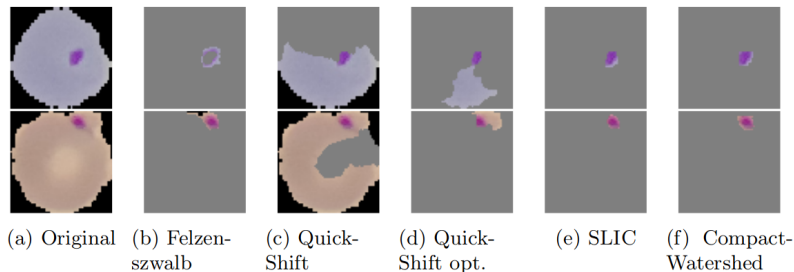
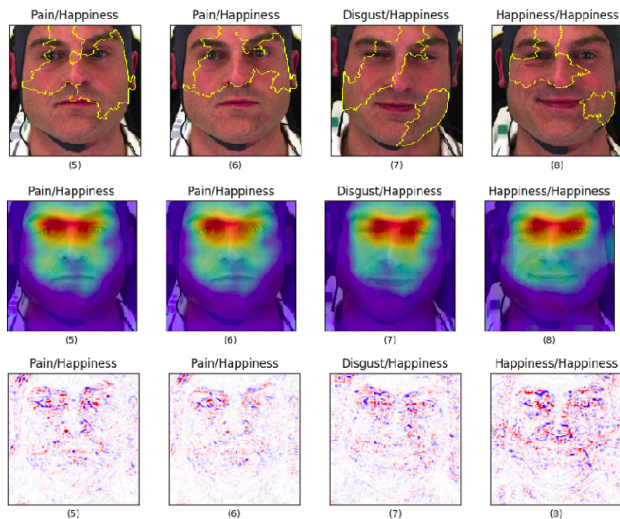


Fig. 4: LIME results for true positive predicted malaria infected cells

Schallner, Rabold, Scholz, Schmid, Effect of Superpixel Aggregation on Explanations in LIME – A Case Study with Biological Data, AIMLA 2019

# Visual Explanations



LIME


CAM

LRP

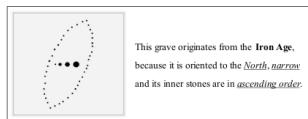
Weitz, Hassan, Schmid, Garbas, Deep-learned faces of pain and emotions: Elucidating the differences of facial expressions with the help of explainable AI methods, tm-Technisches Messen, 2019

# Visual explanations are often not sufficient

- Helpful to recognize overfitting
- Fast communication of information (attention, relevance)
- BUT – visual highlighting is not expressive enough for
  - ▶ spatial relations (the blowhole is **on** a supporting part)
  - ▶ quantification (**all** blowholes are smaller than 1 mm)
  - ▶ feature values (the eyes are **shut** not open)
  - ▶ negation (there is **not** a blowhole but a hairline crack)
  - ▶ recursion (an arbitrary number of objects of increasing size)





# Experimental Findings on Explanations, Joint Performance, and Trust



- (1) **North**  $\wedge$  **Narrow**  $\rightarrow$  **Iron**
- (2) **North**  $\wedge$  **Ascending**  $\rightarrow$  **Iron**
- (3) **Narrow**  $\wedge$  **Ascending**  $\rightarrow$  **Iron**

(4 features, based on Medin & Schaffer, 1978, stimuli pattern for classification learning)

		Verbal explanations	
		available	not available
System error information	available	<i>EG1</i>	<i>EG3</i>
	not available	<i>EG2</i>	<i>CG</i>



(ANOVA,  $n = 190$ ; a priori power analysis for a medium effect size ( $f = .25$ ,  $\alpha = .05$ ,  $1 - \beta = .90$ ) gave a minimum required sample size of 171 participants)

Thaler & Schmid, Explaining Machine Learned Relational Concepts in Visual Domains – Effects of Perceived Accuracy on Joint Performance and Trust, CogSci2021

# Some Observations on Explanations

Tim Miller, AIJ 2019; Tania Lombrozo, TiCS 2006

- There are different possibilities to explain something to someone
  - ▶ verbal (different degrees of detail)
  - ▶ visual (maybe with symbolic annotations)
  - ▶ prototypical examples
  - ▶ contrastive (near miss) example
- There is no one-size fits all (context specificity)
- Explanations can be wrong (*right for the wrong reasons*, Teso & Kersting, AAAI/ACM Conference on AI, Ethics, and Society, 2019)
- Explanations are not always helpful (*Beneficial and Harmful Explanatory Machine Learning*, Ai, Muggleton, . . . , Schmid, MLJ 2021)
- Explanations might lead to unjustified trust
- Explanations can be mutual and extend interactive learning (Schmid & Finzel, KI 2020)
- Explanations are a process

# Contrastive Near-miss Explanations: Structural Alignment

## APPENDIX

Table A1. High- and low-similarity word pairs used in Experiments 1 and 2

Similar pairs		Dissimilar pairs	
Light bulb	Candle	VCR	Lounge chair
Kitten	Cat	Hammock	Horse track
Magazine	Newspaper	Bed	Hockey
Bowl	Mug	Football	Boutique
Phone book	Dictionary	Kite	Painting
Microphone	Stereo speaker	Sculpture	Navy
Piano	Organ	Army	Abacus
Air conditioner	Furnace	Calculator	Escalator
Freezer	Refrigerator	Stairs	Stool
Hammer	Mallet	Broom	Sailboat
Bicycle	Tricycle	Yacht	Missile
Dumpster	Garbage can	Chair	Banana split
Lake	Ocean	Ice cream sundae	Clock
Telephone	CB radio	McDonald's	Couch
Diamond	Ruby	Police car	Burger King
Sponge	Towel	Rocket	Motel
Computer	Typewriter	Hotel	Tape deck
Staple	Paper clip	Watch	Ambulance
Shoe	Sandal	Casino	Mop
Chemistry	Biology	Stove	Hang glider
VCR	Tape deck	Light bulb	Cat
Hammock	Lounge chair	Kitten	Newspaper

(Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science*, 5(3), 152-158.)

# Contrastive Near-miss Explanations: Relational Learning

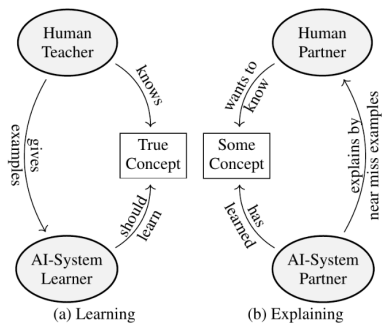


Fig. 2: Duality of learning and explaining




Fig. 5: The positive and negative example structures for the Winston arches domain.

(Rabold, Siebers, Schmid, MLJ, to appear)

# Contrastive Near-miss Explanations: Relational Learning

## Local explanation

A local explanation for a positive example  $P$  is a ground clause  $C\theta$  where  $C \in T$  such that  $P = head(C\theta)$  and  $T \models body(C\theta)$ .

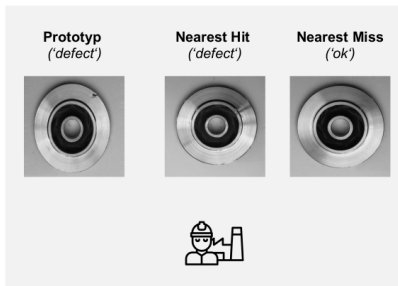

## Near Miss Explanation

Given a local explanation  $C\theta$  and a minimally changed clause  $C'$  with substitution  $\theta'$ , we call  $C'\theta'$  a near miss explanation and  $\Delta head(C'\theta')$  a near miss example if  $T \models body(C'\theta')$ ,  $T \not\models head(C'\theta')$ .

```
struct4
  ↓
arch(A) ← contains(A,X), contains(A,Y), contains(A,Z),
         is_a(X,wedge), is_a(Y,brick), is_a(Z,brick),
         supports(Y,X,A), supports(Z,X,A),
         not_meets(Y,Z,A) meets(b, c, struct4)
```



## Example-based Explainable AI (XAI) Demonstrator




(Fraunhofer IIS CAI, Marvin Herchenbach,2021)

# Explanation of Critical Features by Contrastive Alignment

CogSys Cognitive Companion \* P

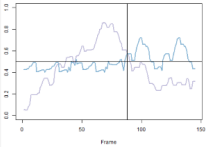
**CGSYS PainFaceComprehender**


**Schmerz** 

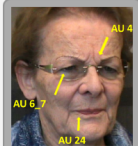
**Warum?**

So unterscheiden sich die Gesichtsausdrücke von Schmerz und Ekel:

**Verlauf der Action Units im Video**



**Schmerz** 

**Ekel** 

# Explanatory Dialogue

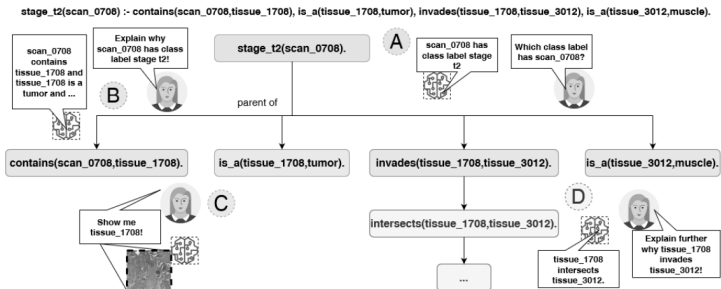


Figure 2: An explanatory tree for `stage_t2(scan_0708)`, that can be queried by the user to get a local explanation why `scan_0708` is labeled as T2 (steps A and B). A dialogue is realized by further requests, either to get more visual explanations in terms of prototypes (step C) or to get more verbal explanations in a drill-down manner (step D).

(Finzel, Tafler, Thaler, Schmid, Multimodal Explanations for User-centric Medical Decision Support Systems, HUMAN.AI @ AAAI 2021; Finzel, Tafler, Scheele, Schmid, Explanation as a process: user-centric construction of multi-level and multi-modal explanations, KI 2021)

# Take Away

- Many application domains have requirements which cannot be met by data intensive blackbox approaches of machine learning alone
- Combining deep learning and ILP supports learning of classifiers for image data together with relational explanations
- Mutual explanations and interactive learning allow to integrate expert/common sense knowledge into the learning process resulting in less need for data and allowing to correct erroneous decisions of the learned model
- Explanations are not one size fits all – therefore research should address different explanation modalities, their combination, and strategies to select the most helpful explanations
- Research on explanations and human-in-the loop ML requires interdisciplinary collaboration with psychology and education

Learning without  
thought is labor  
lost Confucius



# Thanks to Team, Cooperation Partners, Funding Agencies



Bundesministerium  
für Bildung  
und Forschung


TraMeExCo (ML-3)

**DFG**

Deutsche  
Forschungsgemeinschaft

Dare2Del  
(SPP 1921)

PainFaceReader

 **Fraunhofer**  
Projektgruppe IIS  
Comprehensible AI

**bidt** Bayerisches  
Forschungsinstitut für  
Digitale Transformation



Bettina Finzel (XI-ML), Johannes Rabold (Neural-symbolic ML), Mark Gromowski (Logik + Learning), Durgesh Nandini (Interactive ILP), Sarem Seitz (Statistical ML), Gesina Schwalbe (Continental), Deniz Neufeld (Bosch), Sebastian Bruckert (Datev), Dr. Stephan Scheele (embedded AI), Ines Rieger, Dominik Seuß, Jaspar Pahl, Teena Hassan (IIS), Oliver May (IIS Dresden)

Stephen Muggleton (Imperial  
College)

Comelia Niessen (FAU)

Stephan Lautenbacher (Uni BA)

Andreas Holzinger (TU Graz)

Torben Weiss (Uni Duisburg-  
Essen)



**Forschungsgruppe  
Elementarinformatik**