

End-to-End Differentiable Proving

Tim Rocktäschel



Whiteson Research Lab, University of Oxford

<http://rockt.github.com> Twitter: @_rockt tim.rocktaschel@cs.ox.ac.uk


Logic and Learning Workshop at The Alan Turing Institute

January 12, 2018

Joint Work With




Sebastian Riedel
University College London



Pasquale Minervini
University College London



Thomas Demeester
Ghent University



Sameer Singh
University of California, Irvine



What vegetable is on the plate?

Neural Net: broccoli

Ground Truth: broccoli



What color are the shoes on the person's feet ?

Neural Net: brown

Ground Truth: brown



How many school busses are there?

Neural Net: 2

Ground Truth: 2



What sport is this?

Neural Net: baseball

Ground Truth: baseball



What is on top of the refrigerator?

Neural Net: magnets

Ground Truth: cereal



What uniform is she wearing?

Neural Net: shorts

Ground Truth: girl scout



What is the table number?

Neural Net: 4

Ground Truth: 40



What are people sitting under in the back?

Neural Net: bench

Ground Truth: tent

Monet \leftrightarrow Photos



Monet \rightarrow photo




photo \rightarrow Monet

Zebras \leftrightarrow Horses



zebra \rightarrow horse



horse \rightarrow zebra

Summer \leftrightarrow Winter



summer \rightarrow winter




winter \rightarrow summer



Photograph

Monet



Van Gogh



Cezanne




Ukiyo-e


Monet ↔ Photos




Monet → photo




Zebras ↔ Horses



zebra → horse




Summer ↔ Winter



summer → winter




a) Chemical Representation of the Synthesis Plan



b) Search Tree Representation

Root (Target)



A = {1} B = {2,6} C = {3,6}

D = {4,5,6} E = {8,9} F = {7,8}




Monet \leftrightarrow Photos



Zebras \leftrightarrow Horses



Summer \leftrightarrow Winter



→ winter



on



Terminal
solved state

$C = \{3, 6\}$

$F = \{7, 8\}$




Monet ↔ Photos



Zebras ↔ Horses



Summer ↔ Winter

Translate from **GERMAN** (detected) ▼Translate into **ENGLISH** ▼


Die Polizei in den USA darf sich wieder schwere Ausrüstung und Waffen beim Militär besorgen. Das hat US-Präsident Donald Trump entschieden und so eine Anordnung seines Vorgängers Barack Obama aufgehoben, nach der es dem Verteidigungsministerium verboten war, die Polizei mit Granatwerfern, gepanzerten Fahrzeugen, Bajonetten, großkalibrigen Waffen und Munition auszurüsten.

Mit der Maßnahme soll sichergestellt werden, dass die Polizei die lebensrettende Ausrüstung bekomme, die sie brauche, um ihren Job zu machen, sagte US-Justizminister Jeff Sessions.

The police in the USA are allowed to get heavy equipment and weapons from the military again. This was decided by US President Donald Trump, who overturned an order from his predecessor Barack Obama, according to which the Department of Defense was banned from equipping the police with grenade launchers, armoured vehicles, bayonets, large-calibre weapons and ammunition.

The measure is designed to ensure that the police get the life-saving equipment they need to do their job, US Attorney General Jeff Sessions said.


Monet  Photos



Transla

Die Polizei in den USA darf sich keine Waffen beim Militär besorgen. Trump entschieden und so ein Barack Obama aufgehoben, r Verteidigungsministerium ver Granatwerfern, gepanzerten Fahrzeugen, Bajonetten, großkalibrigen Waffen und Munition auszurüsten.

Mit der Maßnahme soll sichergestellt werden, dass die Polizei die lebensrettende Ausrüstung bekomme, die sie brauche, um ihren Job zu machen, sagte US-Justizminister Jeff Sessions.



00014
PLAYER ONE

01
GAME

5
LIVES

...ence was banned from equipping the police with grenades, launchers, armoured vehicles, bayonets, large-calibre weapons and ammunition.

The measure is designed to ensure that the police get the life-saving equipment they need to do their job, US Attorney General Jeff Sessions said.








XKCD, 17th May 2017

Data & Explanations


- Rules
- (Partial) Programs
- Natural Language



XKCD, 17th May 2017

Data & Explanations

- Rules
- (Partial) Programs
- Natural Language




Answers & Explanations

- Rules
- Programs
- Natural Language
- Plans

XKCD, 17th May 2017

- Rules

- (Partial) Programs
- Natural Language



- Rules

- Programs
- Natural Language
- Plans


XKCD, 17th May 2017

Data & Explanations

- Rules
- (Partial) Programs
- Natural Language

Answers & Explanations

- Rules
- Programs
- Natural Language
- Plans



XKCD, 17th May 2017

Lecture Notes

PROLOG AND NATURAL-LANGUAGE ANALYSIS

Fernando C.N. Pereira
and
Stuart M. Shieber

CSLI

CENTER FOR THE STUDY
OF LANGUAGE
AND INFORMATION

goal problem.

rule 1
if not turn_over and
battery_bad
then problem is battery cf 100.


rule 2
if lights_weak
then battery_bad cf 50.

rule 3
if radio_weak
then battery_bad cf 50.

rule 4
if turn_over and
smell_gas
then problem is flooded cf 80.

rule 5
if turn_over and
gas_gauge is empty
then problem is out_of_gas cf 90.

rule 6
if turn_over and
gas_gauge is low
then problem is out_of_gas cf 30.



```
goal problem.
```

```
rule 1
  if not turn_over and
    battery_bad
  then problem is battery cf 100.
```

```
rule 2
  if lights_weak
  then battery_bad cf 50.
```

```
rule 3
  if radio_weak
  then battery_bad cf 50.
```


```
rule 4
  if turn_over and
    smell_gas
  then problem is flooded cf 80.
```

```
rule 5
  if turn_over and
    gas_gauge is empty
  then problem is out_of_gas cf 90.
```

```
rule 6
  if turn_over and
    gas_gauge is low
  then problem is out_of_gas cf 30.
```

Expert Systems

- No/little training data
- Interpretable



```
goal problem.
```

```
rule 1
  if not turn_over and
    battery_bad
  then problem is battery cf 100.
```


```
rule 2
  if lights_weak
  then battery_bad cf 50.
```

```
rule 3
  if radio_weak
  then battery_bad cf 50.
```

```
rule 4
  if turn_over and
    smell_gas
  then problem is flooded cf 80.
```

```
rule 5
  if turn_over and
    gas_gauge is empty
  then problem is out_of_gas cf 90.
```

```
rule 6
  if turn_over and
    gas_gauge is low
  then problem is out_of_gas cf 30.
```



Expert Systems

- No/little training data
- Interpretable
- Rules manually defined
- No generalization

goal problem.

```
rule 1
  if not turn_over and
    battery_bad
  then problem is battery cf 100.
```


```
rule 2
  if lights_weak
  then battery_bad cf 50.
```

```
rule 3
  if radio_weak
  then battery_bad cf 50.
```

```
rule 4
  if turn_over and
    smell_gas
  then problem is flooded cf 80.
```

```
rule 5
  if turn_over and
    gas_gauge is empty
  then problem is out_of_gas cf 90.
```

```
rule 6
  if turn_over and
    gas_gauge is low
  then problem is out_of_gas cf 30.
```



Expert Systems

- No/little training data
- Interpretable
- Rules manually defined
- No generalization

goal problem.

```
rule 1
  if not turn_over and
    battery_bad
  then problem is battery cf 100.
```


```
rule 2
  if lights_weak
  then battery_bad cf 50.
```

```
rule 3
  if radio_weak
  then battery_bad cf 50.
```

```
rule 4
  if turn_over and
    smell_gas
  then problem is flooded cf 80.
```

```
rule 5
  if turn_over and
    gas_gauge is empty
  then problem is out_of_gas cf 90.
```

```
rule 6
  if turn_over and
    gas_gauge is low
  then problem is out_of_gas cf 30.
```



Expert Systems

- No/little training data
- Interpretable
- Rules manually defined
- No generalization



Neural Networks

- Trained end-to-end
- Strong generalization

goal problem.

```
rule 1
  if not turn_over and
    battery_bad
  then problem is battery cf 100.
```


```
rule 2
  if lights_weak
  then battery_bad cf 50.
```

```
rule 3
  if radio_weak
  then battery_bad cf 50.
```

```
rule 4
  if turn_over and
    smell_gas
  then problem is flooded cf 80.
```

```
rule 5
  if turn_over and
    gas_gauge is empty
  then problem is out_of_gas cf 90.
```

```
rule 6
  if turn_over and
    gas_gauge is low
  then problem is out_of_gas cf 30.
```



Expert Systems

- No/little training data
- Interpretable
- Rules manually defined
- No generalization



Neural Networks

- Need a lot of training data
- Not interpretable
- Trained end-to-end
- Strong generalization

goal problem.

```
rule 1
  if not turn_over and
    battery_bad
  then problem is battery cf 100.
```

```
rule 2
  if lights_weak
  then battery_bad cf 50.
```

```
rule 3
  if radio_weak
  then battery_bad cf 50.
```


```
rule 4
  if turn_over and
    smell_gas
  then problem is flooded cf 80.
```

```
rule 5
  if turn_over and
    gas_gauge is empty
  then problem is out_of_gas cf 90.
```

```
rule 6
  if turn_over and
    gas_gauge is low
  then problem is out_of_gas cf 30.
```

Expert Systems

- No/little training data
- Interpretable



Neural Networks

- Trained end-to-end
- Strong generalization

Machine Learning & Logic

- Fuzzy Logic ([Zadeh, 1965](#))

Machine Learning & Logic

- Fuzzy Logic ([Zadeh, 1965](#))
- Probabilistic Logic Programming, e.g.,

Machine Learning & Logic

- Fuzzy Logic ([Zadeh, 1965](#))
- Probabilistic Logic Programming, e.g.,
 - IBAL ([Pfeffer, 2001](#)), BLOG ([Milch et al., 2005](#)), Markov Logic Networks ([Richardson and Domingos, 2006](#)), ProbLog ([De Raedt et al., 2007](#)) . . .

Machine Learning & Logic

- Fuzzy Logic ([Zadeh, 1965](#))
- Probabilistic Logic Programming, e.g.,
 - IBAL ([Pfeffer, 2001](#)), BLOG ([Milch et al., 2005](#)), Markov Logic Networks ([Richardson and Domingos, 2006](#)), ProbLog ([De Raedt et al., 2007](#)) . . .
- Inductive Logic Programming, e.g.,

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...
 - Statistical Predicate Invention (Kok and Domingos, 2007)

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...
 - Statistical Predicate Invention (Kok and Domingos, 2007)
- Neural-symbolic Connectionism

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...
 - Statistical Predicate Invention (Kok and Domingos, 2007)
- Neural-symbolic Connectionism
 - Propositional rules: EBL-ANN (Shavlik and Towell, 1989), KBANN (Towell and Shavlik, 1994), C-LIP (d'Avila Garcez and Zaverucha, 1999)

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...
 - Statistical Predicate Invention (Kok and Domingos, 2007)
- Neural-symbolic Connectionism
 - Propositional rules: EBL-ANN (Shavlik and Towell, 1989), KBANN (Towell and Shavlik, 1994), C-LIP (d'Avila Garcez and Zaverucha, 1999)
 - First-order inference (no training of symbol representations): Unification Neural Networks (Hölldobler, 1990; Komendantskaya, 2011), SHRUTI (Shastri, 1992), Neural Prolog (Ding, 1995), CLIP++ (Franca et al., 2014), Lifted Relational Networks (Sourek et al., 2015)

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...
 - Statistical Predicate Invention (Kok and Domingos, 2007)
- Neural-symbolic Connectionism
 - Propositional rules: EBL-ANN (Shavlik and Towell, 1989), KBANN (Towell and Shavlik, 1994), C-LIP (d'Avila Garcez and Zaverucha, 1999)
 - First-order inference (no training of symbol representations): Unification Neural Networks (Hölldobler, 1990; Komendantskaya, 2011), SHRUTI (Shastri, 1992), Neural Prolog (Ding, 1995), CLIP++ (Franca et al., 2014), Lifted Relational Networks (Sourek et al., 2015)
- Recent: Logic Tensor Networks (Serafini and d'Avila Garcez, 2016), TensorLog (Cohen, 2016), Differentiable Inductive Logic (Evans and Grefenstette, 2017)

Machine Learning & Logic

- Fuzzy Logic (Zadeh, 1965)
- Probabilistic Logic Programming, e.g.,
 - IBAL (Pfeffer, 2001), BLOG (Milch et al., 2005), Markov Logic Networks (Richardson and Domingos, 2006), ProbLog (De Raedt et al., 2007) ...
- Inductive Logic Programming, e.g.,
 - Plotkin (1970), Shapiro (1991), Muggleton (1991), De Raedt (1999) ...
 - Statistical Predicate Invention (Kok and Domingos, 2007)
- Neural-symbolic Connectionism
 - Propositional rules: EBL-ANN (Shavlik and Towell, 1989), KBANN (Towell and Shavlik, 1994), C-LIP (d'Avila Garcez and Zaverucha, 1999)
 - First-order inference (no training of symbol representations): Unification Neural Networks (Hölldobler, 1990; Komendantskaya, 2011), SHRUTI (Shastri, 1992), Neural Prolog (Ding, 1995), CLIP++ (Franca et al., 2014), Lifted Relational Networks (Sourek et al., 2015)
- Recent: Logic Tensor Networks (Serafini and d'Avila Garcez, 2016), TensorLog (Cohen, 2016), Differentiable Inductive Logic (Evans and Grefenstette, 2017)

For overviews see Besold et al. (2017) and d'Avila Garcez et al. (2012)

Outline

- 1** Link prediction & symbolic vs. neural representations

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow ([Rocktäschel et al., 2015](#))

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow (Rocktäschel et al., 2015)
 - Fast but restricted (Demeester et al., 2016)

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow (Rocktäschel et al., 2015)
 - Fast but restricted (Demeester et al., 2016)
 - Model-agnostic and fast (Minervini et al., 2017)

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow (Rocktäschel et al., 2015)
 - Fast but restricted (Demeester et al., 2016)
 - Model-agnostic and fast (Minervini et al., 2017)
- 3 End-to-end differentiable proving (Rocktäschel and Riedel, 2017)

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow (Rocktäschel et al., 2015)
 - Fast but restricted (Demeester et al., 2016)
 - Model-agnostic and fast (Minervini et al., 2017)
- 3 End-to-end differentiable proving (Rocktäschel and Riedel, 2017)
 - Explicit multi-hop reasoning using neural networks

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow (Rocktäschel et al., 2015)
 - Fast but restricted (Demeester et al., 2016)
 - Model-agnostic and fast (Minervini et al., 2017)
- 3 End-to-end differentiable proving (Rocktäschel and Riedel, 2017)
 - Explicit multi-hop reasoning using neural networks
 - Inducing rules using gradient descent

Outline

- 1 Link prediction & symbolic vs. neural representations
- 2 Regularize neural representations using logical rules
 - Model-agnostic but slow (Rocktäschel et al., 2015)
 - Fast but restricted (Demeester et al., 2016)
 - Model-agnostic and fast (Minervini et al., 2017)
- 3 End-to-end differentiable proving (Rocktäschel and Riedel, 2017)
 - Explicit multi-hop reasoning using neural networks
 - Inducing rules using gradient descent
- 4 Outlook & Summary

Notation

- **Constant:** `HOMER`, `BART`, `LISA` etc. (lowercase)

Notation

- **Constant:** HOMER, BART, LISA etc. (lowercase)
- **Variable:** X, Y etc. (uppercase, universally quantified)

Notation

- **Constant:** HOMER, BART, LISA etc. (lowercase)
- **Variable:** X, Y etc. (uppercase, universally quantified)
- **Term:** constant or variable

Restricted to function-free terms in this talk

Notation

- **Constant:** HOMER, BART, LISA etc. (lowercase)
- **Variable:** X, Y etc. (uppercase, universally quantified)
- **Term:** constant or variable
Restricted to function-free terms in this talk
- **Predicate:** fatherOf, parentOf etc.
function from terms to a Boolean

Notation

- **Constant:** HOMER, BART, LISA etc. (lowercase)
- **Variable:** X, Y etc. (uppercase, universally quantified)
- **Term:** constant or variable
Restricted to function-free terms in this talk
- **Predicate:** fatherOf, parentOf etc.
function from terms to a Boolean
- **Atom:** predicate and terms, e.g., parentOf(X, BART)

Notation

- **Constant:** HOMER, BART, LISA etc. (lowercase)
- **Variable:** X, Y etc. (uppercase, universally quantified)
- **Term:** constant or variable
Restricted to function-free terms in this talk
- **Predicate:** fatherOf, parentOf etc.
function from terms to a Boolean
- **Atom:** predicate and terms, e.g., parentOf(X, BART)
- **Literal:** atom or negated or atom, e.g.,
not parentOf(BART, LISA)

Notation

- **Constant:** `HOMER`, `BART`, `LISA` etc. (lowercase)
- **Variable:** `X`, `Y` etc. (uppercase, universally quantified)
- **Term:** constant or variable
Restricted to function-free terms in this talk
- **Predicate:** `fatherOf`, `parentOf` etc.
function from terms to a Boolean
- **Atom:** predicate and terms, e.g., `parentOf(X, BART)`
- **Literal:** atom or negated or atom, e.g.,
not `parentOf(BART, LISA)`
- **Rule:** `head :- body.`
`head`: atom
`body`: (possibly empty) list of literals representing conjunction
Restricted to Horn clauses in this talk

Notation

- **Constant:** `HOMER`, `BART`, `LISA` etc. (lowercase)
- **Variable:** `X`, `Y` etc. (uppercase, universally quantified)
- **Term:** constant or variable
Restricted to function-free terms in this talk
- **Predicate:** `fatherOf`, `parentOf` etc.
function from terms to a Boolean
- **Atom:** predicate and terms, e.g., `parentOf(X, BART)`
- **Literal:** atom or negated or atom, e.g.,
not `parentOf(BART, LISA)`
- **Rule:** `head :- body.`
`head`: atom
`body`: (possibly empty) list of literals representing conjunction
Restricted to Horn clauses in this talk
- **Fact:** ground rule (no free variables) with empty body, e.g.,
`parentOf(HOMER, BART).`

Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

- `placeOfBirth` attribute is missing for 71% of people!

Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

- `placeOfBirth` attribute is missing for 71% of people!
- Commonsense knowledge often not stated explicitly

Link Prediction


Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

- `placeOfBirth` attribute is missing for 71% of people!
- Commonsense knowledge often not stated explicitly
- Weak logical relationships that can be used for inferring facts

Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

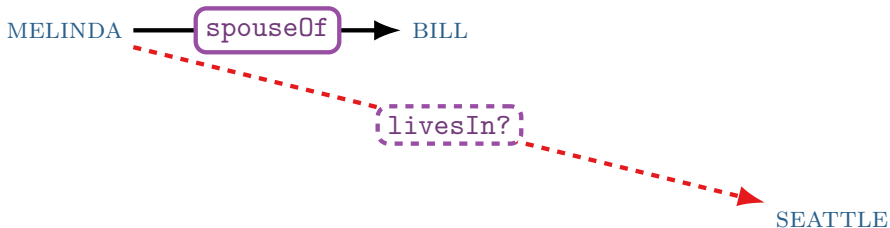
- `placeOfBirth` attribute is missing for 71% of people!
- Commonsense knowledge often not stated explicitly
- Weak logical relationships that can be used for inferring facts



Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

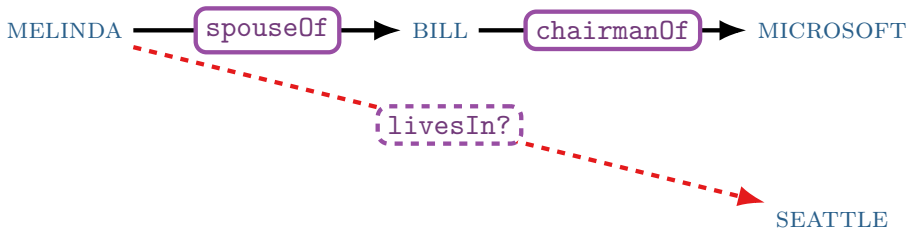
- `placeOfBirth` attribute is missing for 71% of people!
- Commonsense knowledge often not stated explicitly
- Weak logical relationships that can be used for inferring facts



Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

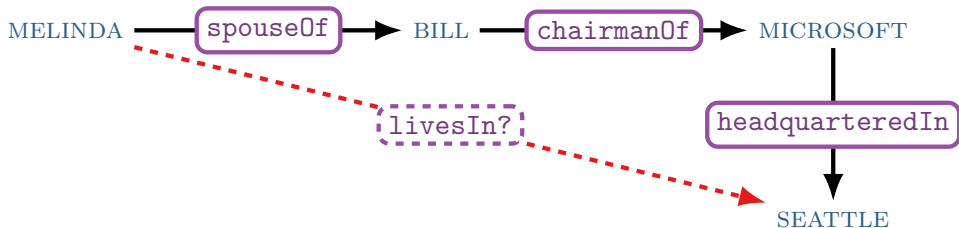
- `placeOfBirth` attribute is missing for 71% of people!
- Commonsense knowledge often not stated explicitly
- Weak logical relationships that can be used for inferring facts



Link Prediction

Real world knowledge bases (like Freebase, DBPedia, YAGO, etc.) are incomplete!

- `placeOfBirth` attribute is missing for 71% of people!
- Commonsense knowledge often not stated explicitly
- Weak logical relationships that can be used for inferring facts



Symbolic Representations

- Symbols (constants and predicates) do not share any information:
`grandpa0f` \neq `grandfather0f`

Symbolic Representations

- Symbols (constants and predicates) do not share any information:
 $\text{grandpaOf} \neq \text{grandfatherOf}$
- No notion of similarity: $\text{APPLE} \sim \text{ORANGE}$, $\text{professorAt} \sim \text{lecturerAt}$

Symbolic Representations

- Symbols (constants and predicates) do not share any information:
`grandpaOf` \neq `grandfatherOf`
- No notion of similarity: `APPLE` \sim `ORANGE`, `professorAt` \sim `lecturerAt`
- No generalization beyond what can be symbolically inferred:
`isFruit`(`APPLE`), `APPLE` \sim `ORANGE`, `isFruit`(`ORANGE`)?

Symbolic Representations

- Symbols (constants and predicates) do not share any information:
`grandpaOf` \neq `grandfatherOf`
- No notion of similarity: `APPLE` \sim `ORANGE`, `professorAt` \sim `lecturerAt`
- No generalization beyond what can be symbolically inferred:
`isFruit`(`APPLE`), `APPLE` \sim `ORGANGE`, `isFruit`(`ORANGE`)?
- Hard to work with language, vision and other modalities
‘‘is a film based on the novel of the same name by’’(`X`, `Y`)

Symbolic Representations

- Symbols (constants and predicates) do not share any information:
 $\text{grandpaOf} \neq \text{grandfatherOf}$
- No notion of similarity: $\text{APPLE} \sim \text{ORANGE}$, $\text{professorAt} \sim \text{lecturerAt}$
- No generalization beyond what can be symbolically inferred:
 $\text{isFruit}(\text{APPLE})$, $\text{APPLE} \sim \text{ORANGE}$, $\text{isFruit}(\text{ORANGE})$?
- Hard to work with language, vision and other modalities
“is a film based on the novel of the same name by”(X, Y)
- But... leads to powerful inference mechanisms and proofs for predictions:
 $\text{fatherOf}(\text{ABE}, \text{HOMER})$. $\text{parentOf}(\text{HOMER}, \text{LISA})$. $\text{parentOf}(\text{HOMER}, \text{BART})$.
 $\text{grandfatherOf}(\text{X}, \text{Y}) \text{ :- } \text{fatherOf}(\text{X}, \text{Z}), \text{parentOf}(\text{Z}, \text{Y})$.
 $\text{grandfatherOf}(\text{ABE}, \text{Q})? \quad \{\text{Q}/\text{LISA}\}, \{\text{Q}/\text{BART}\}$

Symbolic Representations

- Symbols (constants and predicates) do not share any information:
`grandpaOf` \neq `grandfatherOf`
- No notion of similarity: `APPLE` \sim `ORANGE`, `professorAt` \sim `lecturerAt`
- No generalization beyond what can be symbolically inferred:
`isFruit`(`APPLE`), `APPLE` \sim `ORANGE`, `isFruit`(`ORANGE`)?
- Hard to work with language, vision and other modalities
‘‘is a film based on the novel of the same name by’’(`X`, `Y`)
- But... leads to powerful inference mechanisms and proofs for predictions:
`fatherOf`(`ABE`, `HOMER`). `parentOf`(`HOMER`, `LISA`). `parentOf`(`HOMER`, `BART`).
`grandfatherOf`(`X`, `Y`) :- `fatherOf`(`X`, `Z`), `parentOf`(`Z`, `Y`).
`grandfatherOf`(`ABE`, `Q`)? $\{Q/LISA\}, \{Q/BART\}$
- Fairly easy to debug and trivial to incorporate domain knowledge:
Show to domain expert and let her change/add rules and facts

Neural Representations

- Lower-dimensional fixed-length vector representations of symbols (predicates and constants):

$$\mathbf{v}_{\text{APPLE}}, \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{fatherOf}}, \dots \in \mathbb{R}^k$$

Neural Representations

- Lower-dimensional fixed-length vector representations of symbols (predicates and constants):

$$\mathbf{v}_{\text{APPLE}}, \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{fatherOf}}, \dots \in \mathbb{R}^k$$

- Can capture similarity and even semantic hierarchy of symbols:

$$\mathbf{v}_{\text{grandpaOf}} = \mathbf{v}_{\text{grandfatherOf}},$$

$$\mathbf{v}_{\text{APPLE}} \sim \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{APPLE}} < \mathbf{v}_{\text{FRUIT}}$$

Neural Representations

- Lower-dimensional fixed-length vector representations of symbols (predicates and constants):

$$\mathbf{v}_{\text{APPLE}}, \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{fatherOf}}, \dots \in \mathbb{R}^k$$

- Can capture similarity and even semantic hierarchy of symbols:

$$\mathbf{v}_{\text{grandpaOf}} = \mathbf{v}_{\text{grandfatherOf}},$$

$$\mathbf{v}_{\text{APPLE}} \sim \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{APPLE}} < \mathbf{v}_{\text{FRUIT}}$$

- Can be trained from raw task data (e.g. facts in a knowledge base)

Neural Representations

- Lower-dimensional fixed-length vector representations of symbols (predicates and constants):

$$\mathbf{v}_{\text{APPLE}}, \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{fatherOf}}, \dots \in \mathbb{R}^k$$

- Can capture similarity and even semantic hierarchy of symbols:

$$\mathbf{v}_{\text{grandpaOf}} = \mathbf{v}_{\text{grandfatherOf}},$$

$$\mathbf{v}_{\text{APPLE}} \sim \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{APPLE}} < \mathbf{v}_{\text{FRUIT}}$$

- Can be trained from raw task data (e.g. facts in a knowledge base)
- Can be compositional

$$\mathbf{v}_{\text{'is the father of'}} = \text{RNN}_{\theta}(\mathbf{v}_{\text{is}}, \mathbf{v}_{\text{the}}, \mathbf{v}_{\text{father}}, \mathbf{v}_{\text{of}})$$

Neural Representations

- Lower-dimensional fixed-length vector representations of symbols (predicates and constants):

$$\mathbf{v}_{\text{APPLE}}, \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{fatherOf}}, \dots \in \mathbb{R}^k$$

- Can capture similarity and even semantic hierarchy of symbols:

$$\mathbf{v}_{\text{grandpaOf}} = \mathbf{v}_{\text{grandfatherOf}},$$

$$\mathbf{v}_{\text{APPLE}} \sim \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{APPLE}} < \mathbf{v}_{\text{FRUIT}}$$

- Can be trained from raw task data (e.g. facts in a knowledge base)
- Can be compositional

$$\mathbf{v}_{\text{‘‘is the father of’’}} = \text{RNN}_{\theta}(\mathbf{v}_{\text{is}}, \mathbf{v}_{\text{the}}, \mathbf{v}_{\text{father}}, \mathbf{v}_{\text{of}})$$

- But... need large amount of training data

Neural Representations

- Lower-dimensional fixed-length vector representations of symbols (predicates and constants):

$$\mathbf{v}_{\text{APPLE}}, \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{fatherOf}}, \dots \in \mathbb{R}^k$$

- Can capture similarity and even semantic hierarchy of symbols:

$$\mathbf{v}_{\text{grandpaOf}} = \mathbf{v}_{\text{grandfatherOf}},$$

$$\mathbf{v}_{\text{APPLE}} \sim \mathbf{v}_{\text{ORANGE}}, \mathbf{v}_{\text{APPLE}} < \mathbf{v}_{\text{FRUIT}}$$

- Can be trained from raw task data (e.g. facts in a knowledge base)
- Can be compositional

$$\mathbf{v}_{\text{“is the father of”}} = \text{RNN}_{\theta}(\mathbf{v}_{\text{is}}, \mathbf{v}_{\text{the}}, \mathbf{v}_{\text{father}}, \mathbf{v}_{\text{of}})$$

- But... need large amount of training data
- No direct way of incorporating prior knowledge

$$\mathbf{v}_{\text{grandfatherOf}}(X, Y) \text{ :- } \mathbf{v}_{\text{fatherOf}}(X, Z), \mathbf{v}_{\text{parentOf}}(Z, Y).$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \mathbf{v}_s^{\top} (\mathbf{v}_i \odot \mathbf{v}_j) \\ &= \sum_k \mathbf{v}_{sk} \mathbf{v}_{ik} \mathbf{v}_{jk} \end{aligned}$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \mathbf{v}_s^{\top} (\mathbf{v}_i \odot \mathbf{v}_j) \\ &= \sum_k \mathbf{v}_{sk} \mathbf{v}_{ik} \mathbf{v}_{jk} \end{aligned}$$

Complex (Trouillon et al., 2016)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{C}^k$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \mathbf{v}_s^{\top} (\mathbf{v}_i \odot \mathbf{v}_j) \\ &= \sum_k \mathbf{v}_{sk} \mathbf{v}_{ik} \mathbf{v}_{jk} \end{aligned}$$

Complex (Trouillon et al., 2016)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{C}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \\ &\text{real}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \\ &+ \text{real}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &+ \text{imag}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &- \text{imag}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \end{aligned}$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \mathbf{v}_s^{\top} (\mathbf{v}_i \odot \mathbf{v}_j) \\ &= \sum_k \mathbf{v}_{sk} \mathbf{v}_{ik} \mathbf{v}_{jk} \end{aligned}$$

Complex (Trouillon et al., 2016)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{C}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \\ &\text{real}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \\ &+ \text{real}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &+ \text{imag}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &- \text{imag}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \end{aligned}$$

Training Loss

$$\mathcal{L} = \sum_{r_s(e_i, e_j), y \in \mathcal{T}} -y \log(\sigma(f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j))) - (1 - y) \log(1 - \sigma(f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j)))$$

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \mathbf{v}_s^{\top} (\mathbf{v}_i \odot \mathbf{v}_j) \\ &= \sum_k \mathbf{v}_{sk} \mathbf{v}_{ik} \mathbf{v}_{jk} \end{aligned}$$

Complex (Trouillon et al., 2016)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{C}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \\ &\text{real}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \\ &+ \text{real}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &+ \text{imag}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &- \text{imag}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \end{aligned}$$

Training Loss

$$\mathcal{L} = \sum_{r_s(e_i, e_j), y \in \mathcal{T}} -y \log(\sigma(f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j))) - (1 - y) \log(1 - \sigma(f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j)))$$

- Learn $\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j$ from data

State-of-the-art Neural Link Prediction

$$\text{livesIn}(\text{MELINDA}, \text{SEATTLE})? = f_{\theta}(\mathbf{v}_{\text{livesIn}}, \mathbf{v}_{\text{MELINDA}}, \mathbf{v}_{\text{SEATTLE}})$$

DistMult (Yang et al., 2015)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{R}^k$$

$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \mathbf{v}_s^{\top} (\mathbf{v}_i \odot \mathbf{v}_j) \\ &= \sum_k \mathbf{v}_{sk} \mathbf{v}_{ik} \mathbf{v}_{jk} \end{aligned}$$

ComplEx (Trouillon et al., 2016)

$$\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j \in \mathbb{C}^k$$


$$\begin{aligned} f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j) &= \\ &\text{real}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \\ &+ \text{real}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &+ \text{imag}(\mathbf{v}_s)^{\top} (\text{real}(\mathbf{v}_i) \odot \text{imag}(\mathbf{v}_j)) \\ &- \text{imag}(\mathbf{v}_s)^{\top} (\text{imag}(\mathbf{v}_i) \odot \text{real}(\mathbf{v}_j)) \end{aligned}$$

Training Loss

$$\mathcal{L} = \sum_{r_s(e_i, e_j), y \in \mathcal{T}} -y \log(\sigma(f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j))) - (1 - y) \log(1 - \sigma(f_{\theta}(\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j)))$$


- Learn $\mathbf{v}_s, \mathbf{v}_i, \mathbf{v}_j$ from data
- Obtain gradients $\nabla_{\mathbf{v}_s} \mathcal{L}, \nabla_{\mathbf{v}_i} \mathcal{L}, \nabla_{\mathbf{v}_j} \mathcal{L}$ by backprop

Regularization by Propositional Logic



Regularization by Propositional Logic


$\text{fatherOf}(X, Y) \text{ :- } \text{parentOf}(X, Y), \neg \text{motherOf}(X, Y)$



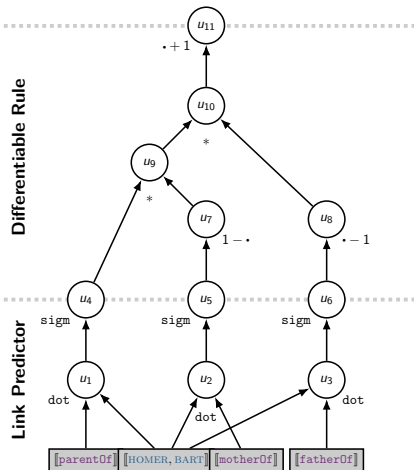
Regularization by Propositional Logic

$\text{fatherOf}(X, Y) :- \text{parentOf}(X, Y), \neg \text{motherOf}(X, Y)$

$$p(F) = \llbracket F \rrbracket = \begin{cases} f_{\theta}(s, i, j) & \text{if } F = s(i, j) \\ 1 - \llbracket A \rrbracket & \text{if } F = \neg A \\ \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \wedge B \\ \llbracket A \rrbracket + \llbracket B \rrbracket - \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \vee B \\ \llbracket B \rrbracket (\llbracket A \rrbracket - 1) + 1 & \text{if } F = A :- B \end{cases}$$



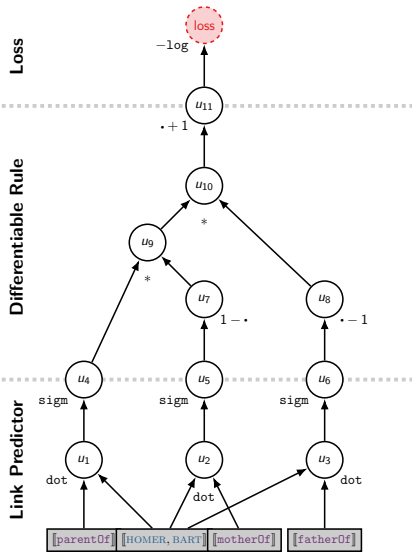
Regularization by Propositional Logic



$\text{fatherOf}(X, Y) :- \text{parentOf}(X, Y), \neg \text{motherOf}(X, Y)$

$$p(F) = \llbracket F \rrbracket = \begin{cases} f_{\theta}(s, i, j) & \text{if } F = s(i, j) \\ 1 - \llbracket A \rrbracket & \text{if } F = \neg A \\ \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \wedge B \\ \llbracket A \rrbracket + \llbracket B \rrbracket - \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \vee B \\ \llbracket B \rrbracket (\llbracket A \rrbracket - 1) + 1 & \text{if } F = A :- B \end{cases}$$

Regularization by Propositional Logic



$\text{fatherOf}(X, Y) :- \text{parentOf}(X, Y), \neg \text{motherOf}(X, Y)$

$$p(F) = \llbracket F \rrbracket = \begin{cases} f_{\theta}(s, i, j) & \text{if } F = s(i, j) \\ 1 - \llbracket A \rrbracket & \text{if } F = \neg A \\ \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \wedge B \\ \llbracket A \rrbracket + \llbracket B \rrbracket - \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \vee B \\ \llbracket B \rrbracket (\llbracket A \rrbracket - 1) + 1 & \text{if } F = A :- B \end{cases}$$

$\mathcal{L}(\llbracket \text{fatherOf}(\text{HOMER}, \text{BART}) :- \text{parentOf}(\text{HOMER}, \text{BART}) \wedge \neg \text{motherOf}(\text{HOMER}, \text{BART}) \rrbracket)$


Regularization by Propositional Logic

$$\mathcal{L}(F) = -\log(\llbracket \forall \mathbf{X}, \mathbf{Y} : F(\mathbf{X}, \mathbf{Y}) \rrbracket) = -\sum_{(e_i, e_j) \in \mathcal{C}^2} \log \llbracket F(e_i, e_j) \rrbracket$$

Loss

Differentiable Rule

Link Predictor




$$\text{fatherOf}(\mathbf{X}, \mathbf{Y}) \text{ :- } \text{parentOf}(\mathbf{X}, \mathbf{Y}), \neg \text{motherOf}(\mathbf{X}, \mathbf{Y})$$

$$p(F) = \llbracket F \rrbracket = \begin{cases} f_{\theta}(s, i, j) & \text{if } F = s(i, j) \\ 1 - \llbracket A \rrbracket & \text{if } F = \neg A \\ \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \wedge B \\ \llbracket A \rrbracket + \llbracket B \rrbracket - \llbracket A \rrbracket \llbracket B \rrbracket & \text{if } F = A \vee B \\ \llbracket B \rrbracket (\llbracket A \rrbracket - 1) + 1 & \text{if } F = A \text{ :- } B \end{cases}$$

$$\mathcal{L}(\llbracket \text{fatherOf}(\text{HOMER}, \text{BART}) \text{ :- } \text{parentOf}(\text{HOMER}, \text{BART}) \wedge \neg \text{motherOf}(\text{HOMER}, \text{BART}) \rrbracket)$$


Zero-shot Learning Results

■ Neural Link Prediction (LP)




Zero-shot Learning Results

■ Neural Link Prediction (LP) ■ Deduction




Zero-shot Learning Results

■ Neural Link Prediction (LP) ■ Deduction ■ Deduction after LP




Zero-shot Learning Results



Zero-shot Learning Results


■ Neural Link Prediction (LP) ■ Deduction ■ Deduction after LP
■ Deduction before LP ■ Regularization



Lifted Regularization by Implications

Every **father** is a **parent**


Every **mother** is a **parent**



Lifted Regularization by Implications



Every **father** is a **parent**

Every **mother** is a **parent**





Lifted Regularization by Implications

Every **father** is a **parent**
Every **mother** is a **parent**



Lifted Regularization by Implications

Every **father** is a **parent**
Every **mother** is a **parent**





Lifted Regularization by Implications

Every **father** is a **parent**

Generalises to similar relations (e.g. **dad**)

Every **mother** is a **parent**

Generalises to similar relations (e.g. **mum**)



Lifted Regularization by Implications

Every **father** is a **parent**

Every **mother** is a **parent**


Every **parent** is a **relative**

Generalises to similar relations (e.g. **dad**)


Generalises to similar relations (e.g. **mum**)

No training facts needed!

Before



After




Lifted Regularization by Implications


Every **father** is a **parent**
Every **mother** is a **parent**
Every **parent** is a **relative**

$$\begin{aligned} \forall X, Y : h(X, Y) &:- b(X, Y) \\ \forall (e_i, e_j) \in \mathcal{C}^2 : \llbracket h \rrbracket^\top \llbracket e_i, e_j \rrbracket &\geq \llbracket b \rrbracket^\top \llbracket e_i, e_j \rrbracket \\ \llbracket h \rrbracket &\geq \llbracket b \rrbracket, \quad \forall (e_i, e_j) \in \mathcal{C}^2 : \llbracket e_i, e_j \rrbracket \in \mathbb{R}_+^k \end{aligned}$$

Before



After




Adversarial Regularization

Clause \mathcal{A} : $h(\mathbf{X}, \mathbf{Y}) \text{ :- } b_1(\mathbf{X}, \mathbf{Z}) \wedge b_2(\mathbf{Z}, \mathbf{Y})$




- Regularization by propositional rules needs grounding – does not scale to large domains!



Adversarial Regularization

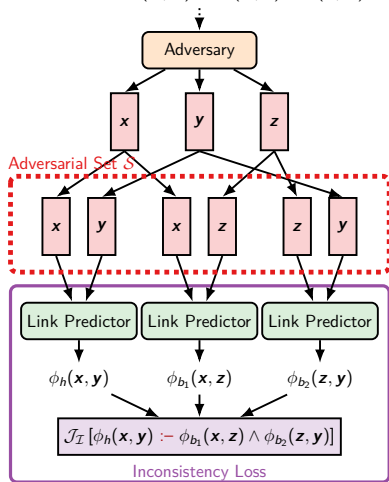
Clause \mathcal{A} : $h(\mathbf{X}, \mathbf{Y}) \text{ :- } b_1(\mathbf{X}, \mathbf{Z}) \wedge b_2(\mathbf{Z}, \mathbf{Y})$



- Regularization by propositional rules needs grounding – does not scale to large domains!
- Lifted regularization only supports direct implications

Adversarial Regularization


Clause \mathcal{A} : $h(\mathbf{X}, \mathbf{Y}) \text{ :- } b_1(\mathbf{X}, \mathbf{Z}) \wedge b_2(\mathbf{Z}, \mathbf{Y})$



- Regularization by propositional rules needs grounding – does not scale to large domains!
- Lifted regularization only supports direct implications
- Idea: let grounding be generated by an adversary and optimize minimax game...

Adversarial Regularization

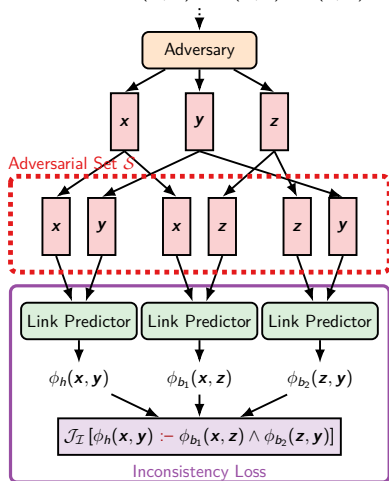
Clause \mathcal{A} : $h(\mathbf{X}, \mathbf{Y}) \text{ :- } b_1(\mathbf{X}, \mathbf{Z}) \wedge b_2(\mathbf{Z}, \mathbf{Y})$



- Regularization by propositional rules needs grounding – does not scale to large domains!
- Lifted regularization only supports direct implications
- Idea: let grounding be generated by an adversary and optimize minimax game...
- Adversary finds maximally violating grounding for a given rule

Adversarial Regularization

Clause \mathcal{A} : $h(\mathbf{X}, \mathbf{Y}) \text{ :- } b_1(\mathbf{X}, \mathbf{Z}) \wedge b_2(\mathbf{Z}, \mathbf{Y})$



- Regularization by propositional rules needs grounding – does not scale to large domains!
- Lifted regularization only supports direct implications
- Idea: let grounding be generated by an adversary and optimize minimax game...
- Adversary finds maximally violating grounding for a given rule
- Neural link predictor attempts to minimize rule violation for given generated groundings

End-to-End Differentiable Prover

- Neural network for proving queries to a knowledge base

End-to-End Differentiable Prover

- Neural network for proving queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols

End-to-End Differentiable Prover

- Neural network for proving queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Learn vector representations of symbols** end-to-end from proof success

End-to-End Differentiable Prover

- Neural network for proving queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Learn vector representations of symbols** end-to-end from proof success
- **Make use of provided rules** in soft proofs

End-to-End Differentiable Prover

- Neural network for proving queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Learn vector representations of symbols** end-to-end from proof success
- **Make use of provided rules** in soft proofs
- **Induce interpretable rules** end-to-end from proof success

Approach



Approach



Yann LeCun

@ylecun

Replying to @NandoDF

sort of like "kernelize" used to be.

10:11 AM - 5 Aug 2016

Approach



Yann LeCun

@ylecun

Replying to @NandoDF

sort of like "kernelize" used to be.

10:11 AM - 5 Aug 2016

Let's **neuralize** Prolog's Backward Chaining using a Radial Basis Function **kernel** for unifying vector representations of symbols!

Prolog's Backward Chaining

Example Knowledge Base:



1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`

Prolog's Backward Chaining

Example Knowledge Base:

1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`

Intuition:

- Backward chaining translates a query into subqueries via rules, e.g.,
`grandfatherOf(ABE, BART)`   `fatherOf(ABE, Z), parentOf(Z, BART)`

Prolog's Backward Chaining

Example Knowledge Base:

1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`

Intuition:

- Backward chaining translates a query into subqueries via rules, e.g.,
`grandfatherOf(ABE, BART)` \rightarrow `3.` \rightarrow `fatherOf(ABE, Z), parentOf(Z, BART)`
- It attempts this for all rules in the knowledge base and thus specifies a depth-first search

Unification

Example Knowledge Base:

1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`

Query

`grandfatherOf ABE BART`

Unification

Example Knowledge Base:

1. fatherOf(ABE, HOMER).
2. parentOf(HOMER, BART).
3. grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).

Query

grandfatherOf ABE BART


1. fatherOf ABE HOMER

Unification

Example Knowledge Base:

1. fatherOf(ABE, HOMER).
2. parentOf(HOMER, BART).
3. grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).


Query



Unification

Example Knowledge Base:



1. fatherOf(ABE, HOMER).
2. parentOf(HOMER, BART).
3. grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).



Unification

Example Knowledge Base:


1. fatherOf(ABE, HOMER).
2. parentOf(HOMER, BART).
3. grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).



Unification

Example Knowledge Base:


1. fatherOf(ABE, HOMER).
2. parentOf(HOMER, BART).
3. grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).



Unification

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Unification

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Unification Failure

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Neural Unification

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Neural Unification

Example Knowledge Base:

1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`



$$\min \left(1.0, \exp \left(\frac{-\| \mathbf{v}_{\text{grandpaOf}} - \mathbf{v}_{\text{grandfatherOf}} \|^2}{2\mu^2} \right) \right)$$

Differentiable Prover

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:

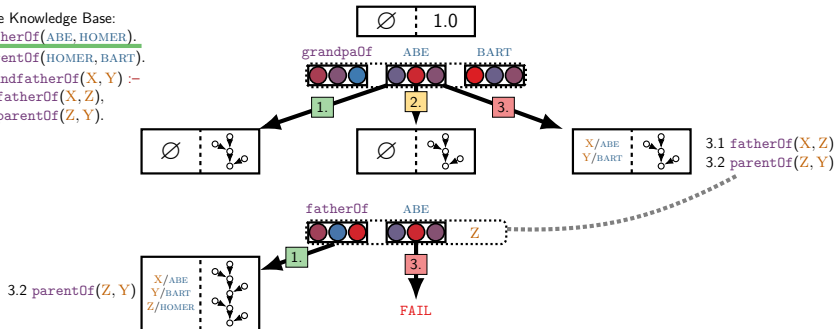
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:

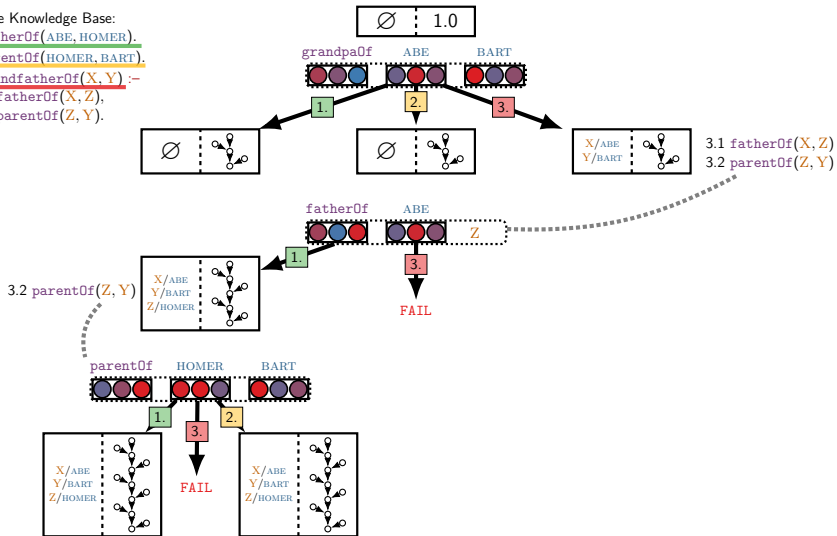
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:

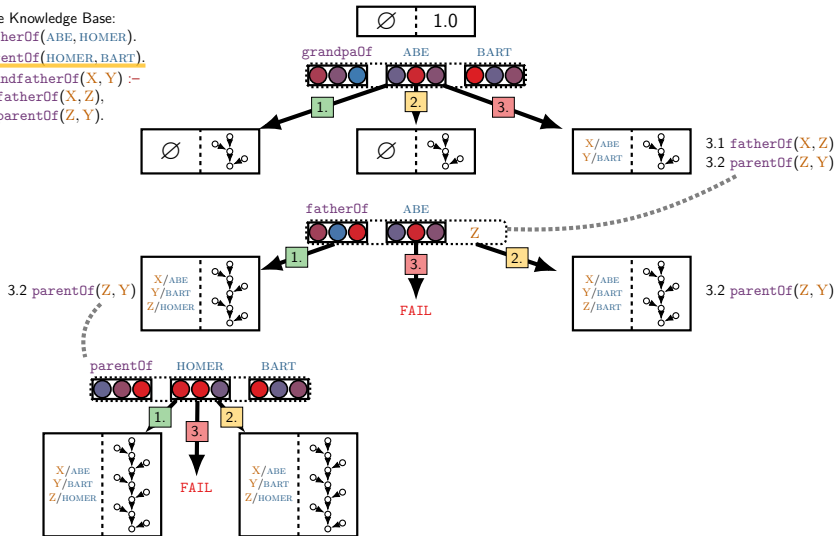
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:

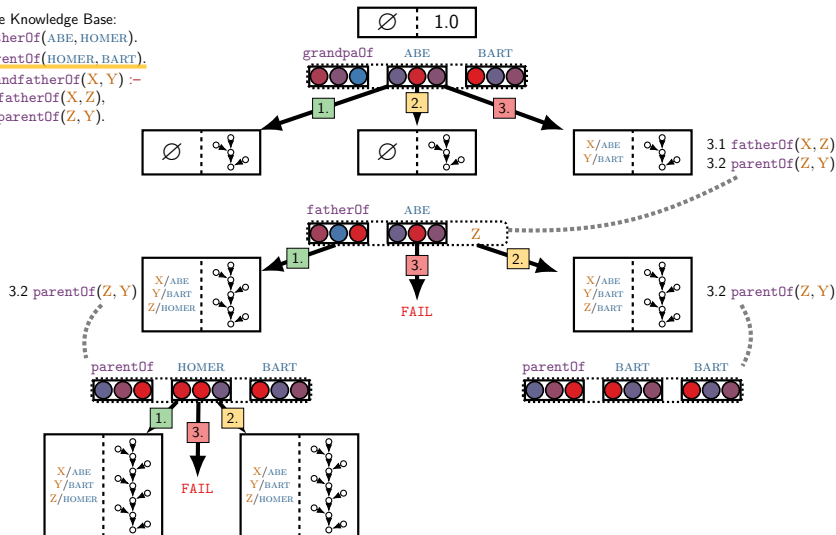
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:

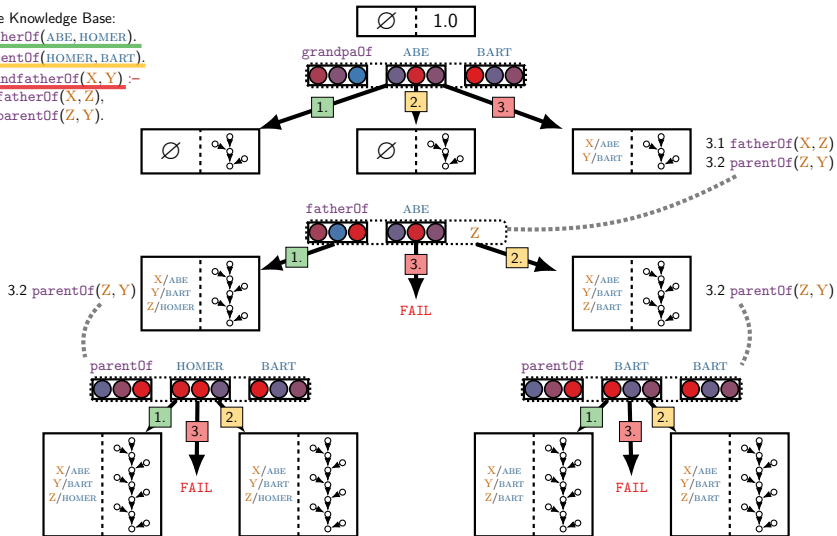
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
`fatherOf(X, Z),`
`parentOf(Z, Y).`



Differentiable Prover

Example Knowledge Base:

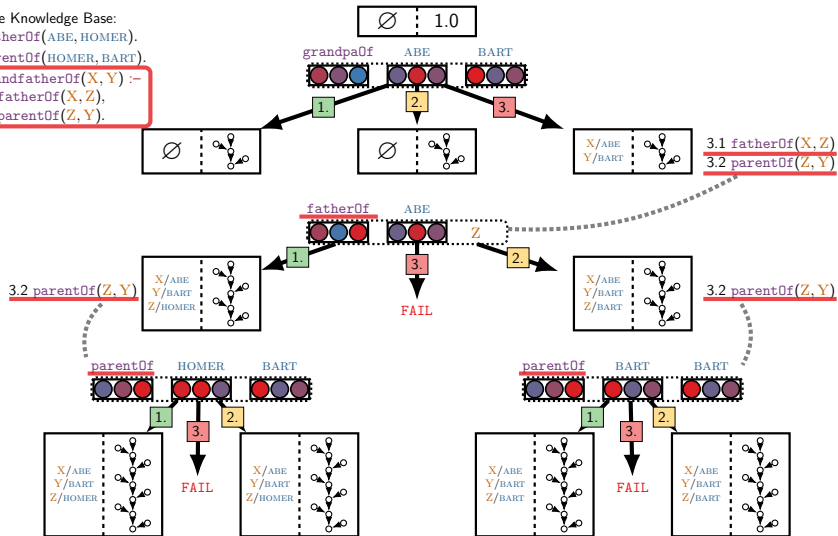
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-`
 `fatherOf(X, Z),`
 `parentOf(Z, Y).`



Neural Program Induction

Example Knowledge Base:

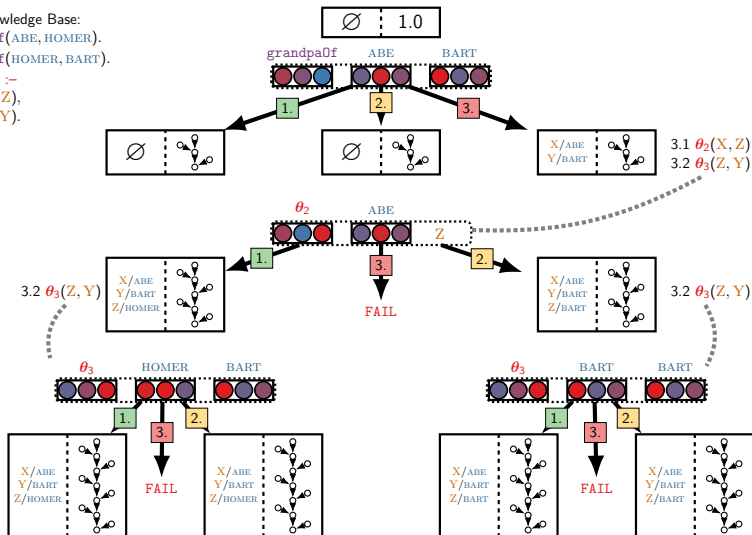
1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. `grandfatherOf(X, Y) :-
 fatherOf(X, Z),
 parentOf(Z, Y).`




Neural Program Induction

Example Knowledge Base:


1. `fatherOf(ABE, HOMER).`
2. `parentOf(HOMER, BART).`
3. $\theta_1(X, Y) :-$
 $\theta_2(X, Z),$
 $\theta_3(Z, Y).$



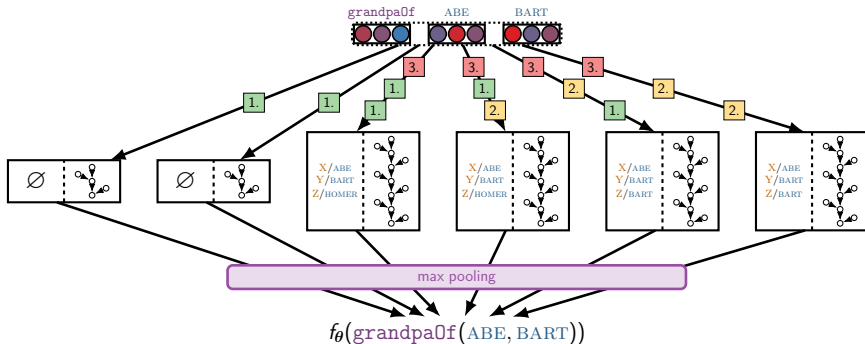
Training Objective



Training Objective

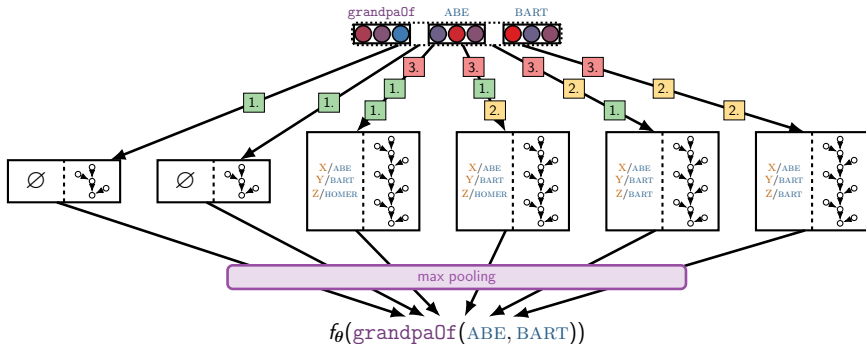


Training Objective



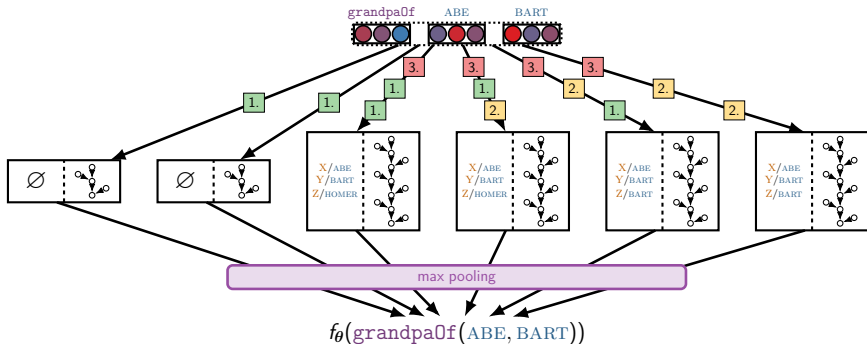
- Loss: negative log-likelihood w.r.t. target proof success

Training Objective




- Loss: negative log-likelihood w.r.t. target proof success
- Trained end-to-end using stochastic gradient descent

Training Objective




- Loss: negative log-likelihood w.r.t. target proof success
- Trained end-to-end using stochastic gradient descent
- Vectors are **learned such that proof success is high for known facts** and low for sampled negative facts

Calculation on GPU




Q

Calculation on GPU




Calculation on GPU

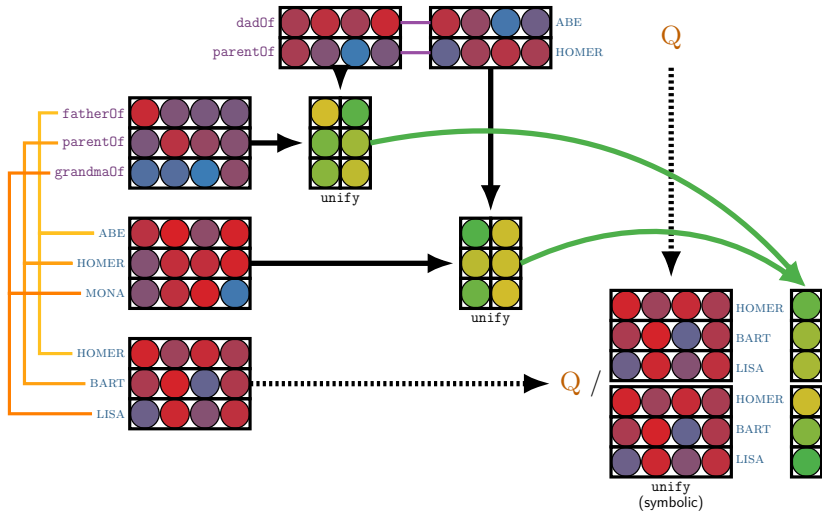


Q

Calculation on GPU





Calculation on GPU





Experiments

Benchmark Knowledge Bases: **Kinship**, **Nations**, **UMLS** (Kok and Domingos, 2007), and **Countries** (Bouchard et al., 2015)





Experiments

Benchmark Knowledge Bases: **Kinship**, **Nations**, **UMLS** (Kok and Domingos, 2007), and **Countries** (Bouchard et al., 2015)





Experiments

Benchmark Knowledge Bases: **Kinship**, **Nations**, **UMLS** (Kok and Domingos, 2007), and **Countries** (Bouchard et al., 2015)



Experiments

Benchmark Knowledge Bases: **Kinship**, **Nations**, **UMLS** (Kok and Domingos, 2007), and **Countries** (Bouchard et al., 2015)



Details

- Models implemented in TensorFlow

Details

- Models implemented in TensorFlow
 - ComplEx** Neural link prediction model by [Trouillon et al. \(2016\)](#)

Details

- Models implemented in TensorFlow

- Complex** Neural link prediction model by [Trouillon et al. \(2016\)](#)

- Prover** End-to-end differentiable prover

Details

- Models implemented in TensorFlow
 - Complex** Neural link prediction model by [Trouillon et al. \(2016\)](#)
 - Prover** End-to-end differentiable prover
 - Prover λ** Same, but representations trained with Complex as auxiliary task

Details

■ Models implemented in TensorFlow

Complex Neural link prediction model by Trouillon et al. (2016)

Prover End-to-end differentiable prover

Prover λ Same, but representations trained with Complex as auxiliary task

■ Rule Templates:

Kinship, Nations & UMLS

20 #1(X, Y) :- #2(X, Y).

20 #1(X, Y) :- #2(Y, X).

20 #1(X, Y) :- #2(X, Z), #3(Z, Y).

Countries S1

3 #1(X, Y) :- #1(Y, X).

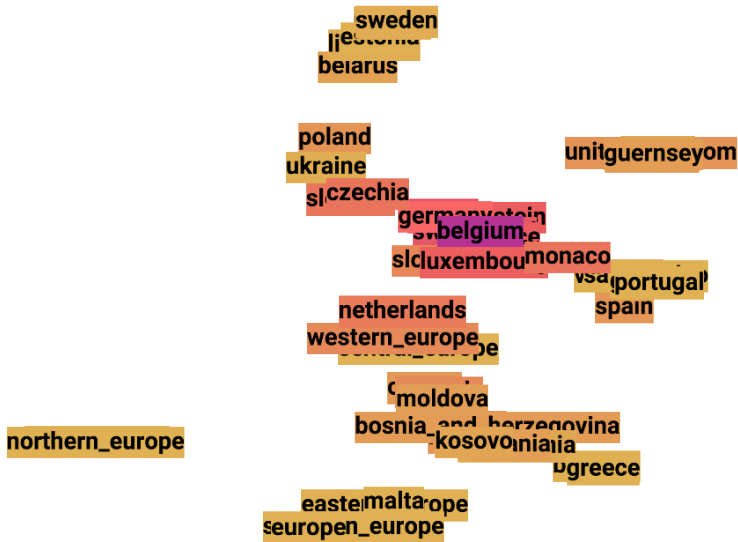
3 #1(X, Y) :- #2(X, Z), #2(Z, Y).

Countries S2


3 #1(X, Y) :- #2(X, Z), #3(Z, Y).

Countries S3


3 #1(X, Y) :- #2(X, Z), #3(Z, W), #4(W, Y).




Results



Results




Results



Examples of Induced Rules

Corpus		Induced rules and their confidence
Countries	S1	0.90 <code>locatedIn(X,Y) :- locatedIn(X,Z), locatedIn(Z,Y).</code>
	S2	0.63 <code>locatedIn(X,Y) :- neighborOf(X,Z), locatedIn(Z,Y).</code>
	S3	0.32 <code>locatedIn(X,Y) :- neighborOf(X,Z), neighborOf(Z,W), locatedIn(W,Y).</code>
Nations		0.68 <code>blockpositionindex(X,Y) :- blockpositionindex(Y,X).</code>
		0.46 <code>expeldiplomats(X,Y) :- negativebehavior(X,Y).</code>
		0.38 <code>negativecomm(X,Y) :- commonbloc0(X,Y).</code>
		0.38 <code>intergovorgs3(X,Y) :- intergovorgs(Y,X).</code>
UMLS		0.88 <code>interacts_with(X,Y) :- interacts_with(X,Z), interacts_with(Z,Y).</code>
		0.77 <code>isa(X,Y) :- isa(X,Z), isa(Z,Y).</code>
		0.71 <code>derivative_of(X,Y) :- derivative_of(X,Z), derivative_of(Z,Y).</code>


Outlook




Outlook

Question

My patient is not responding after three days of codeine treatment.
What could have happened?




Outlook

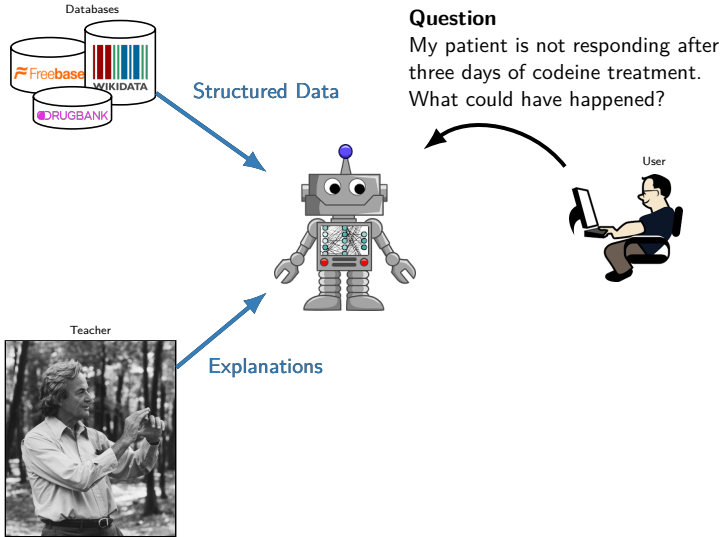


Question

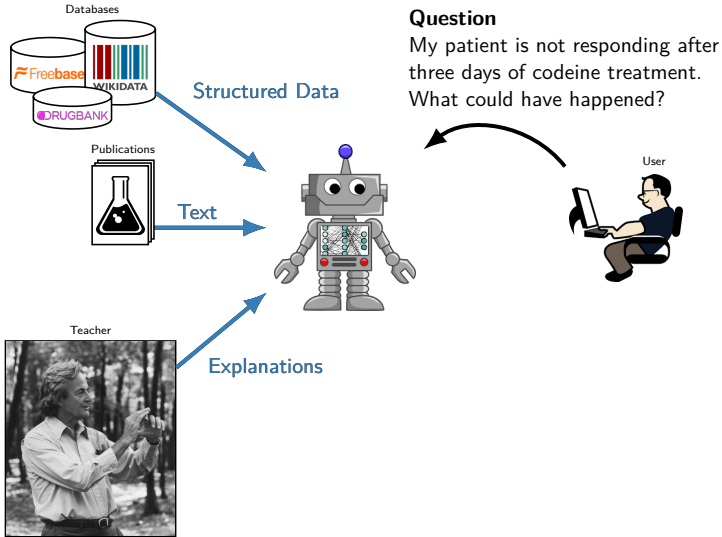
My patient is not responding after three days of codeine treatment.
What could have happened?



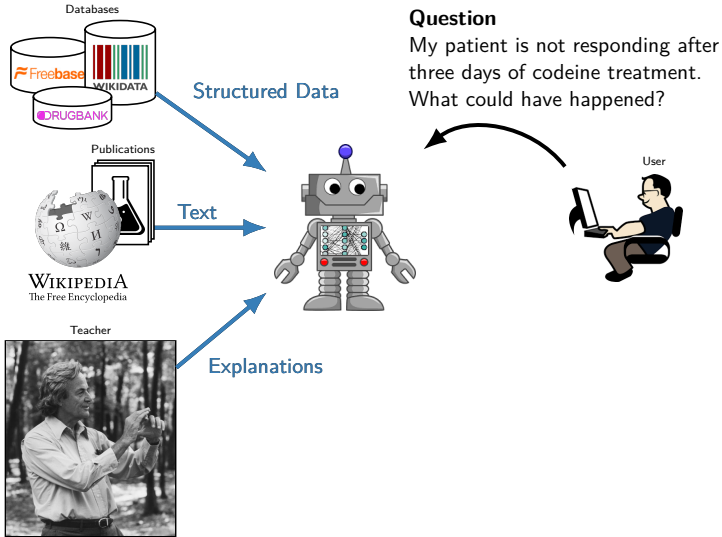
Outlook



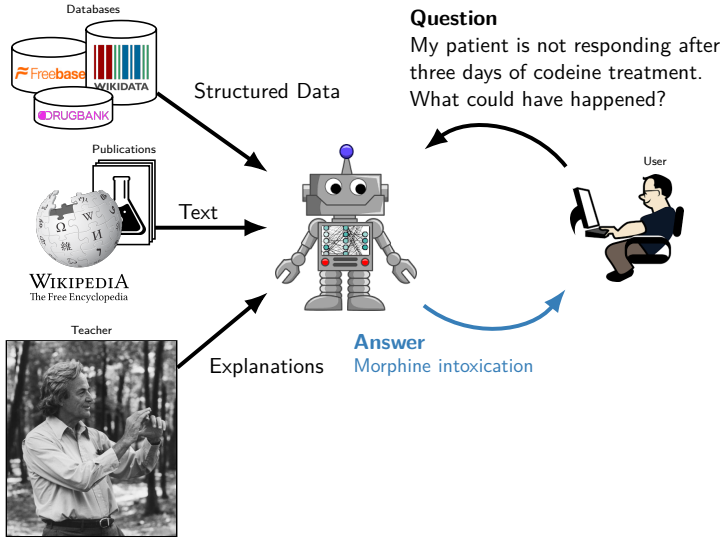
Outlook



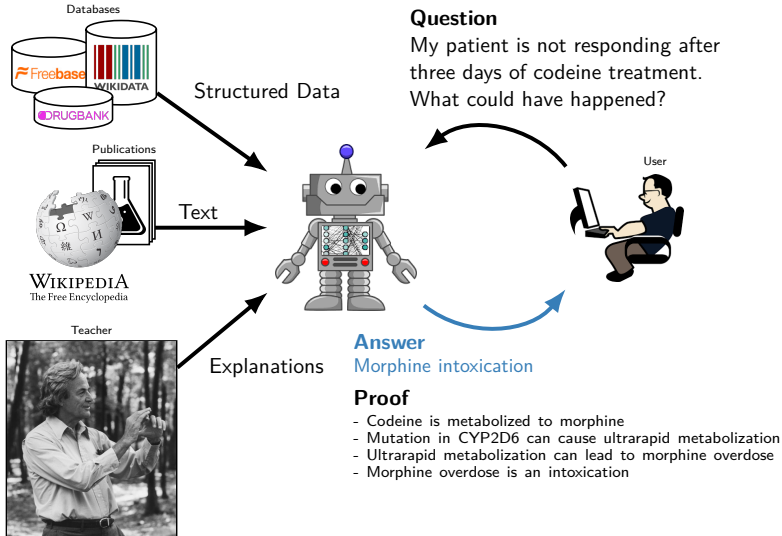
Outlook



Outlook



Outlook



Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**
- **Learns vector representations of symbols** from data via gradient descent

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**
- **Learns vector representations of symbols** from data via gradient descent
- **Induces interpretable rules** from data via gradient descent

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**
- **Learns vector representations of symbols** from data via gradient descent
- **Induces interpretable rules** from data via gradient descent
- Various computational optimizations: batch proving, tree pruning etc.

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**
- **Learns vector representations of symbols** from data via gradient descent
- **Induces interpretable rules** from data via gradient descent
- Various computational optimizations: batch proving, tree pruning etc.
- Future research:

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**
- **Learns vector representations of symbols** from data via gradient descent
- **Induces interpretable rules** from data via gradient descent
- Various computational optimizations: batch proving, tree pruning etc.
- Future research:
 - **Scaling up** to larger knowledge bases

Summary

- We proposed various ways of **regularizing vector representations of symbols using rules**
- We used Prolog's backward chaining as recipe for recursively constructing a neural network to prove queries to a knowledge base
- Proof success differentiable w.r.t. vector representations of symbols
- **Symbolic rule application but neural unification**
- **Learns vector representations of symbols** from data via gradient descent
- **Induces interpretable rules** from data via gradient descent
- Various computational optimizations: batch proving, tree pruning etc.
- Future research:
 - **Scaling up** to larger knowledge bases
 - **Connecting to RNNs** for proving with natural language statements

Thank you!

<http://rockt.github.com>

tim.rocktaschel@cs.ox.ac.uk

Twitter: @_rockt

References I

- T. R. Besold, A. S. d'Avila Garcez, S. Bader, H. Bowman, P. M. Domingos, P. Hitzler, K. Kühnberger, L. C. Lamb, D. Lowd, P. M. V. Lima, L. de Penning, G. Pinkas, H. Poon, and G. Zaverucha. Neural-symbolic learning and reasoning: A survey and interpretation. *CoRR*, abs/1711.03902, 2017. URL <http://arxiv.org/abs/1711.03902>.
- G. Bouchard, S. Singh, and T. Trouillon. On approximate reasoning capabilities of low-rank vector spaces. In *Proceedings of the 2015 AAAI Spring Symposium on Knowledge Representation and Reasoning (KRR): Integrating Symbolic and Neural Approaches*, 2015.
- W. W. Cohen. Tensorlog: A differentiable deductive database. *CoRR*, abs/1605.06523, 2016. URL <http://arxiv.org/abs/1605.06523>.
- R. Das, A. Neelakantan, D. Belanger, and A. McCallum. Chains of reasoning over entities, relations, and text using recurrent neural networks. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, 2017. URL <http://arxiv.org/abs/1607.01426>.
- A. S. d'Avila Garcez and G. Zaverucha. The connectionist inductive learning and logic programming system. *Appl. Intell.*, 11(1):59–77, 1999. doi: 10.1023/A:1008328630915. URL <http://dx.doi.org/10.1023/A:1008328630915>.
- A. S. d'Avila Garcez, K. Broda, and D. M. Gabbay. *Neural-symbolic learning systems: foundations and applications*. Springer Science & Business Media, 2012.
- T. Demeester, **T. Rocktäschel**, and S. Riedel. Lifted rule injection for relation embeddings. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 1389–1399, 2016. URL <http://aclweb.org/anthology/D/D16/D16-1146.pdf>.
- L. Ding. Neural prolog-the concepts, construction and mechanism. In *Systems, Man and Cybernetics, 1995. Intelligent Systems for the 21st Century., IEEE International Conference on*, volume 4, pages 3603–3608. IEEE, 1995.
- R. Evans and E. Grefenstette. Learning explanatory rules from noisy data. *CoRR*, abs/1711.04574, 2017. URL <http://arxiv.org/abs/1711.04574>.
- S. Hölldobler. A structured connectionist unification algorithm. In *Proceedings of the 8th National Conference on Artificial Intelligence. Boston, Massachusetts, July 29 - August 3, 1990, 2 Volumes.*, pages 587–593, 1990. URL <http://www.aaai.org/Library/AAAI/1990/aaai90-088.php>.

References II

- S. Kok and P. M. Domingos. Statistical predicate invention. In *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20-24, 2007*, pages 433–440, 2007. doi: 10.1145/1273496.1273551. URL <http://doi.acm.org/10.1145/1273496.1273551>.
- E. Komendantskaya. Unification neural networks: unification by error-correction learning. *Logic Journal of the IGPL*, 19(6):821–847, 2011. doi: 10.1093/jigpal/jzq012. URL <http://dx.doi.org/10.1093/jigpal/jzq012>.
- P. Minervini, T. Demeester, **T. Rocktäschel**, and S. Riedel. Adversarial sets for regularised neural link predictors. In *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
- T. Rocktäschel** and S. Riedel. End-to-end differentiable proving. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 3791–3803, 2017. URL <http://papers.nips.cc/paper/6969-end-to-end-differentiable-proving>.
- T. Rocktäschel**, S. Singh, and S. Riedel. Injecting logical background knowledge into embeddings for relation extraction. In *NAACL HLT 2015, The 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Denver, Colorado, USA, May 31 - June 5, 2015*, pages 1119–1129, 2015. URL <http://aclweb.org/anthology/N/N15/N15-1118.pdf>.
- L. Serafini and A. S. d’Avila Garcez. Logic tensor networks: Deep learning and logical reasoning from data and knowledge. In *Proceedings of the 11th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy’16) co-located with the Joint Multi-Conference on Human-Level Artificial Intelligence (HLAI 2016), New York City, NY, USA, July 16-17, 2016.*, 2016. URL http://ceur-ws.org/Vol-1768/NESY16_paper3.pdf.
- L. Shastri. Neurally motivated constraints on the working memory capacity of a production system for parallel processing: Implications of a connectionist model based on temporal synchrony. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society: July 29 to August 1, 1992, Cognitive Science Program, Indiana University, Bloomington*, volume 14, page 159. Psychology Press, 1992.
- J. W. Shavlik and G. G. Towell. An approach to combining explanation-based and neural learning algorithms. *Connection Science*, 1(3): 231–253, 1989.

References III

- G. Sourek, V. Aschenbrenner, F. Zelezný, and O. Kuzelka. Lifted relational neural networks. In *Proceedings of the NIPS Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches co-located with the 29th Annual Conference on Neural Information Processing Systems (NIPS 2015)*, Montreal, Canada, December 11-12, 2015., 2015. URL http://ceur-ws.org/Vol-1583/CoCoNIPS_2015_paper_7.pdf.
- G. G. Towell and J. W. Shavlik. Knowledge-based artificial neural networks. *Artif. Intell.*, 70(1-2):119–165, 1994. doi: 10.1016/0004-3702(94)90105-8. URL [http://dx.doi.org/10.1016/0004-3702\(94\)90105-8](http://dx.doi.org/10.1016/0004-3702(94)90105-8).
- T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pages 2071–2080, 2016. URL <http://jmlr.org/proceedings/papers/v48/trouillon16.html>.
- B. Yang, W. Yih, X. He, J. Gao, and L. Deng. Embedding entities and relations for learning and inference in knowledge bases. In *International Conference on Learning Representations (ICLR)*, 2015. URL <http://arxiv.org/abs/1412.6575>.