Why Isn't Relational Learning Taking Over the World?

David Poole

Department of Computer Science,
University of British Columbia

January 22, 2026

Relational Learning

Relational learning is learning models that make (probabilistic) predictions about entities (things, objects, including events):

- their properties
- relations among them
- existence
- identity

Relational Learning

Relational learning is learning models that make (probabilistic) predictions about entities (things, objects, including events):

- their properties
- relations among them
- existence
- identity

Also statistical relational AI,
relational probabilistic models,
logic learning

Relational Models

"The mind is a neural computer, fitted by natural selection with combinatorial algorithms for causal and probabilistic reasoning about plants, animals, objects, and people.

How the Mind Works, Steven Pinker, 1997

Relational Models

"The mind is a neural computer, fitted by natural selection with combinatorial algorithms for causal and probabilistic reasoning about plants, animals, objects, and people.

How the Mind Works, Steven Pinker, 1997

Relations and Knowledge Graphs

- Example relation:

Patient	Test	Technician	Result	DateTime
54326	353	99807	1	202601220945
54326	353	87601	0	202601250830
...				

Relations and Knowledge Graphs

- Example relation:

Patient	Test	Technician	Result	DateTime
54326	353	99807	1	202601220945
54326	353	87601	0	202601250830

...

- Relational datasets typically contain lots of identifiers: typically arbitrary integers

Relations and Knowledge Graphs

- Example relation:

Patient	Test	Technician	Result	DateTime
54326	353	99807	1	202601220945
54326	353	87601	0	202601250830

...

- Relational datasets typically contain lots of identifiers: typically arbitrary integers
- (subject, verb, object) triples $\longrightarrow$ knowledge graphs.

Relations and Knowledge Graphs

- Example relation:

Patient	Test	Technician	Result	DateTime
54326	353	99807	1	202601220945
54326	353	87601	0	202601250830
...				

- Relational datasets typically contain lots of identifiers: typically arbitrary integers
- (subject, verb, object) triples $\longrightarrow$ knowledge graphs.
- Given a table: use (row, column, value) where row is either a primary key or a reified entity.

Relations and Knowledge Graphs

- Example relation:

Patient	Test	Technician	Result	DateTime
54326	353	99807	1	202601220945
54326	353	87601	0	202601250830
...				

- Relational datasets typically contain lots of identifiers: typically arbitrary integers
- (subject, verb, object) triples $\longrightarrow$ knowledge graphs.
- Given a table: use (row, column, value) where row is either a primary key or a reified entity.
- E.g, (54326353202601220945, patient, 54326) (54326353202601220945, result, 1)

Relations and Knowledge Graphs

- Example relation:

Patient	Test	Technician	Result	DateTime
54326	353	99807	1	202601220945
54326	353	87601	0	202601250830
...				

- Relational datasets typically contain lots of identifiers: typically arbitrary integers
- (subject, verb, object) triples $\longrightarrow$ knowledge graphs.
- Given a table: use (row, column, value) where row is either a primary key or a reified entity.
- E.g, (54326353202601220945, patient, 54326) (54326353202601220945, result, 1)
- Number of triples for reified entity = number of columns

Standard Evaluation Datasets

- FB15k is extracted from Freebase, a predecessor of WikiData
- entities and relations that appeared in more than 100 triples

Standard Evaluation Datasets

- FB15k is extracted from Freebase, a predecessor of WikiData
- entities and relations that appeared in more than 100 triples
- Example triples:

```
(A.S. Livorno Calcio, /soccer/roster,
         Forward (association football))
(Forward (association football), /soccer/position,
         Cambridge United F.C.)
(California, religion, Methodism)
(Ambient music, /music/genre, Portishead (band))
(Marriage, /marriage/spouse, Noel Gallagher)
(Hannah Montana: The Movie,
      film_release_region, Egypt)
```

Reified Entities

- In Wikidata, over 98% of the entities appear as the subject of fewer than 10 triples.
- This includes:
  - Long tail of entities that are just stubs
  - Reified entities.

Reified Entities

- In Wikidata, over 98% of the entities appear as the subject of fewer than 10 triples.
- This includes:
  - Long tail of entities that are just stubs
  - Reified entities.
- Including only entities appearing in many tuples $\longrightarrow$ no reified entities.

Evaluation

How to evaluate a prediction

- Accuracy

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.99999999995% accurate.

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.99999999995% accurate.
- Log loss or cross entropy loss

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.99999999995% accurate.
- Log loss or cross entropy loss: with no negative statements, can't evaluate probabilities.

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.99999999995% accurate.
- Log loss or cross entropy loss: with no negative statements, can't evaluate probabilities.
- Ranking: given test triple (s, v, o), rank object for $(s, v, ?)$. The ranking is the position of o in ranked list.

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.9999999995% accurate.
- Log loss or cross entropy loss: with no negative statements, can't evaluate probabilities.
- Ranking: given test triple (s, v, o), rank object for $(s, v, ?)$. The ranking is the position of o in ranked list.
  - Hit at k

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.99999999995% accurate.
- Log loss or cross entropy loss: with no negative statements, can't evaluate probabilities.
- Ranking: given test triple (s, v, o), rank object for $(s, v, ?)$. The ranking is the position of o in ranked list.
  - Hit at k
  - Mean ranking

Evaluation

How to evaluate a prediction

- Accuracy: To predict whether a random triple is in Wikidata, answering "no" is over 99.99999999995% accurate.
- Log loss or cross entropy loss: with no negative statements, can't evaluate probabilities.
- Ranking: given test triple (s, v, o), rank object for $(s, v, ?)$. The ranking is the position of o in ranked list.
  - Hit at k
  - Mean ranking
  - Mean reciprocal ranking

Evaluation

Problems with ranking:

Evaluation

Problems with ranking:

- Some questions cannot be asked:
  Who is the Pope married to?

Evaluation

Problems with ranking:

- Some questions cannot be asked:
  Who is the Pope married to?
- Asking $(s, v, ?)$ leaks information about the test set.
  Asking who e is married to.

Evaluation

Problems with ranking:

- Some questions cannot be asked:
  Who is the Pope married to?
- Asking $(s, v, ?)$ leaks information about the test set.
  Asking who e is married to.
- Some queries are trivial:
  What positions does Cambridge United F.C. have?

Evaluation

Problems with ranking:

- Some questions cannot be asked:
  Who is the Pope married to?
- Asking $(s, v, ?)$ leaks information about the test set.
  Asking who e is married to.
- Some queries are trivial:
  What positions does Cambridge United F.C. have?
- Some queries are almost impossible:
  Which football team has someone playing *forward*?

Evaluation

Problems with ranking:

- Some questions cannot be asked:
  Who is the Pope married to?
- Asking $(s, v, ?)$ leaks information about the test set.
  Asking who e is married to.
- Some queries are trivial:
  What positions does Cambridge United F.C. have?
- Some queries are almost impossible:
  Which football team has someone playing *forward*?
- Actual probabilities are lost.

Evaluation

Problems with ranking:

- Some questions cannot be asked:
  Who is the Pope married to?
- Asking $(s, v, ?)$ leaks information about the test set.
  Asking who e is married to.
- Some queries are trivial:
  What positions does Cambridge United F.C. have?
- Some queries are almost impossible:
  Which football team has someone playing *forward*?
- Actual probabilities are lost.
- Loses sight of the downstream task.

Looking Forward: missing data

- Ignoring missing data can lead to misleading results.

Looking Forward: missing data

- Ignoring missing data can lead to misleading results.
- Missing data is rarely missing at random:
  - Wikidata contains all Taylor Swift albums released
  - For most recording artists, most albums are missing

Looking Forward: missing data

- Ignoring missing data can lead to misleading results.
- Missing data is rarely missing at random:
  - Wikidata contains all Taylor Swift albums released
  - For most recording artists, most albums are missing
- Open Problem: better ways of handling missing data for relational domains

Looking Forward

- Need better datasets.
  - Most real-world databases are valuable and so secret
  - Use public datasets. Environmental? Scientific data?

Looking Forward

- Need better datasets.
  - Most real-world databases are valuable and so secret
  - Use public datasets. Environmental? Scientific data?
- To made decision we need to be able to infer probabilities:
  - negative information or meta-information ("no more", functional)
  - predict future from past (assume dataset is complete for some cases)

Looking Forward

- Need better datasets.
  - Most real-world databases are valuable and so secret
  - Use public datasets. Environmental? Scientific data?
- To made decision we need to be able to infer probabilities:
  - negative information or meta-information ("no more", functional)
  - predict future from past (assume dataset is complete for some cases)
- Open Problem: Better evaluation for the various types of relational domains and questions

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities
  - an entity or entities not represented (without identifier)

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities
  - an entity or entities not represented (without identifier)
  - no entity

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities
  - an entity or entities not represented (without identifier)
  - no entity
- E.g. asking for birth mother of a person

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities
  - an entity or entities not represented (without identifier)
  - no entity
- E.g. asking for birth mother of a person
- E.g., asking for children of a person

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities
  - an entity or entities not represented (without identifier)
  - no entity
- E.g. asking for birth mother of a person
- E.g., asking for children of a person
- E.g., asking whether the president of South Korea and the prime minister of Canada had a private meeting at the 2025 ASEAN Summit (or 2026 Summit)

Looking Forward: Query Types

- When predicting entities, there are three types:
  - one or more known entities
  - an entity or entities not represented (without identifier)
  - no entity
- E.g. asking for birth mother of a person
- E.g., asking for children of a person
- E.g., asking whether the president of South Korea and the prime minister of Canada had a private meeting at the 2025 ASEAN Summit (or 2026 Summit)
- Open Problem: Determining probabilities for various types of possible answers.

Looking Forward: Aggregation

- Aggregation: prediction depends on variable number of related entities.
  e.g. predicting gender from movies rated.
  Number of related entities can vary from zero to thousands

Looking Forward: Aggregation

- Aggregation: prediction depends on variable number of related entities.
  e.g. predicting gender from movies rated.
  Number of related entities can vary from zero to thousands
- Some models have built-in aggregation
  (e.g., noisy-or, logistic regression, attention)

Looking Forward: Aggregation

- Aggregation: prediction depends on variable number of related entities.
  e.g. predicting gender from movies rated.
  Number of related entities can vary from zero to thousands
- Some models have built-in aggregation
  (e.g., noisy-or, logistic regression, attention)
- Other model use explicit operations
  (e.g., maximum, sum or mean.)

Looking Forward: Aggregation

- Aggregation: prediction depends on variable number of related entities.
  e.g. predicting gender from movies rated.
  Number of related entities can vary from zero to thousands
- Some models have built-in aggregation
  (e.g., noisy-or, logistic regression, attention)
- Other model use explicit operations
  (e.g., maximum, sum or mean.)
- Open Problem: Do related entities provide independent evidence?

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible
- Some ways forward:

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible
- Some ways forward:
  - more types of real-world prediction problems

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible
- Some ways forward:
  - more types of real-world prediction problems
  - learn directly on relations (knowledge hypergraphs)

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible
- Some ways forward:
  - more types of real-world prediction problems
  - learn directly on relations (knowledge hypergraphs)
  - use (public) complete knowledge databases, predict future from past

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible
- Some ways forward:
  - more types of real-world prediction problems
  - learn directly on relations (knowledge hypergraphs)
  - use (public) complete knowledge databases, predict future from past
- Potential to learn from all data in the world.

Conclusions

- Relational learning: model objects and relationships, rather than their manifestations in language or images.
- Why Isn't Relational Learning Taking Over the World?
  - Poor test sets. E.g, contain no reified entities
  - Evaluation: ranking is rarely sensible
- Some ways forward:
  - more types of real-world prediction problems
  - learn directly on relations (knowledge hypergraphs)
  - use (public) complete knowledge databases, predict future from past
- Potential to learn from all data in the world.
- Much more detail in paper.

See also:

Invited talk: "The Essence of Intelligence is Appropriate Action (not thinking, reasoning, learning or language) and other things every student of AI should know"

David Poole and Alan Mackworth

Sunday at 8:30am.