Machine learning with Semantic Web ontologies

Robert Hoehndorf

King Abdullah University of Science and Technology

September 8, 2025

Logic as foundation of AI

Newell & Simon, 1976

A physical symbol system has the necessary and sufficient means for general intelligent action.

- Symbols are physical entities that can encode for our *knowledge* of the world (e.g., for a proposition P)
  - formal and natural languages
- P may not be *explicitly* represented $\Rightarrow$ need for symbol manipulation
  - reasoning, deduction
- main methods: mathematical logic, formal systems

Logic as foundation of AI

- We want our knowledge of the world to affect our decisions:
  - do action A if world believed to satisfy P
  - not: if P is in the knowledge base
- Example:
  - Patient x is allergic to medication m: $all(x, m)$
  - Anybody allergic to m is also allergic to m': $\forall y(all(y, m) \rightarrow all(y, m'))$
  - It's not safe to prescribe medication if allergic: $\forall a, b(all(a, b) \rightarrow \neg safe(a, b))$
  - Is it safe to prescribe m' to x: $safe(x, m')$?

Biology is a knowledge-driven discipline

- micro-theories, knowledge bases/graphs, systems biology models:
  - "judgments" resulting from experiments, supported by evidence
    - `FOXP2 binds CUX1`
    - `TCP1 has-function 'macrophage apoptosis'`
- from > 100 years of experiments
  - can no longer be "learned" from data...
  - ... because data is not available
- biology developed many large, widely used knowledge bases

Ontologies and knowledge graphs

knowledge graphs:

- focus on entities, facts, queries, associations
- open or closed world semantics
- weak semantics (assertional)
- query answering
- potentially very large

OWL ontologies:

- focus on concepts, (complex) axioms, reasoning, consistency
- open world semantics
- model theory
- deduction and proof theory
- usually small

Neuro-symbolic integration in bioinformatics

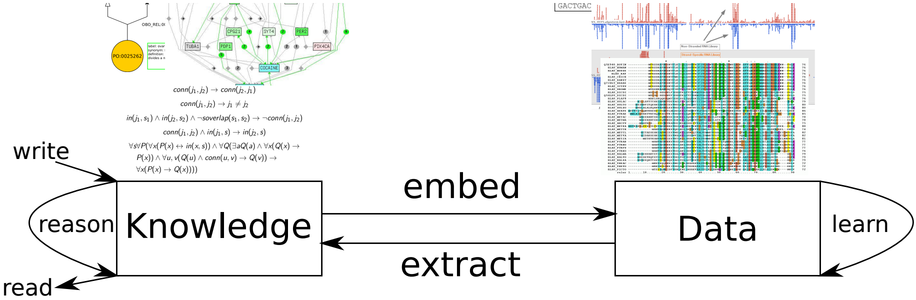

Neuro-symbolic integration in bioinformatics

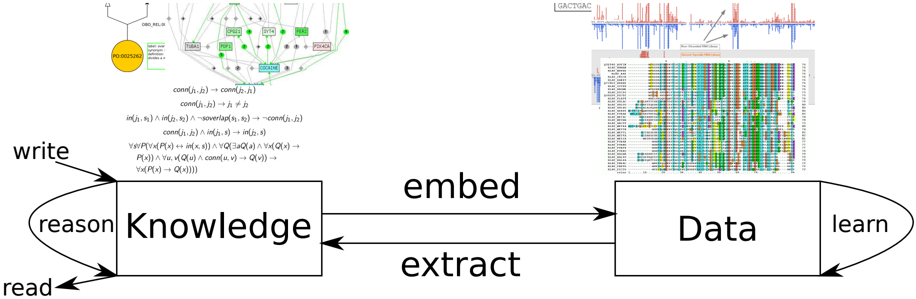

- phenotype
- function
- disease

- genotype
- protein sequence
- expression

Embedding formal knowledge

Embedding

An embedding is a map (morphism) from one mathematical structure X into another structure Y:
$f : X \hookrightarrow Y$
such that X is preserved in Y.

- Y may be more suitable than X for some operations/algorithms.
  - similarity
  - gradients, optimization

Embedding formal knowledge

Embedding

An embedding is a map (morphism) from one mathematical structure X into another structure Y:
$f : X \hookrightarrow Y$
such that X is preserved in Y.

- Y may be more suitable than X for some operations/algorithms.
  - similarity
  - gradients, optimization

For knowledge-enhanced learning, we want to embed formalized *knowledge bases* in $\mathbb{R}^n$. Approaches:

- graph-based, syntactic, model-theoretic

Own work

- 2017: Walking RDF & OWL (graph, entailment)
- 2018: SmuDGE (graph, “semantic” random walk)
- 2018: Onto2Graph (graph projection)
- 2018: Onto2Vec (syntax)
- 2019: OPA2Vec (syntax, natural language)
- 2019: EL Embeddings (models)
- 2022: EL Box Embeddings (models, intersection)
- 2023: FALCON (models, fuzzy logic, theory)
- 2024: CatE (category, graph projection)
- 2024: DELE (deductive closure)

Graph projections

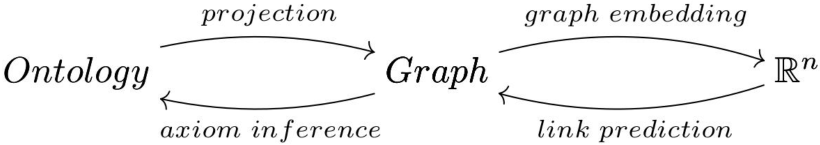

Graph projections

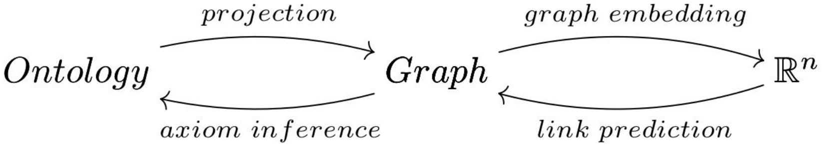

- Knowledge Graph Embeddings, e.g.:
- TransE
- ConvE
- ...

Graph projections

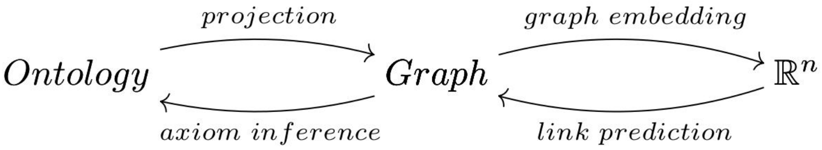

- **DL2Vec**
- **Onto2Graph**
- OWL2Vec*
- RDF syntax trees
- Taxonomy

- Knowledge Graph Embeddings, e.g.:
- TransE
- ConvE
- . . .

Graph projection

Language-based (DL2Vec):

Condition 1	Condition 2	Triple(s)
$A \sqsubseteq QR_0 \ldots QR_m D$ $A \equiv QR_0 \ldots QR_m D$	$D := B_1 \sqcup \ldots \sqcup B_n \mid$ $B_1 \sqcap \ldots \sqcap B_n$	$\langle A, (R_0 ... R_m), B_i \rangle$ for $i \in 1 \ldots n$
$A \sqsubseteq B$		$\langle A, SubClassOf, B \rangle$
$A \equiv B$		$\langle A, EquivalentTo, B \rangle$

Domain-specific (Onto2Graph):

- $X \sqsubseteq Y$: $X \xrightarrow{\text{is-a}} Y$
- $X \sqsubseteq \exists \text{part-of}.Y$: $X \xrightarrow{\text{part-of}} Y$
- $X \sqsubseteq \exists \text{regulates}.Y$: $X \xrightarrow{\text{regulates}} Y$
- $X \sqcap Y \sqsubseteq \bot$: $X \xleftrightarrow{\text{disjoint}} Y$
- $X \equiv Y$: $X \overset{\equiv}{\leftrightarrow} Y$, $\{X, Y\}$

Properties of graph projections

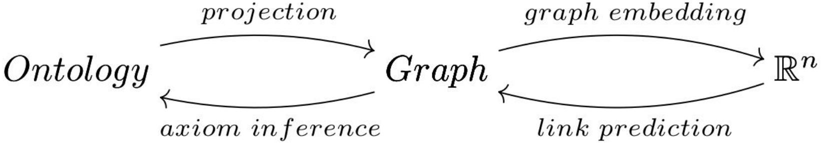

Some graph projections are not embeddings:

Properties of graph projections

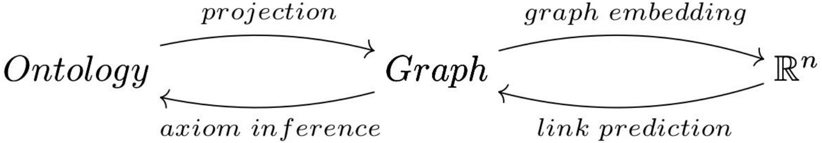

Some graph projections are not embeddings:

- injectivity: $C \sqsubseteq \exists R.D$ and $C \sqsubseteq \forall R.D$

Properties of graph projections

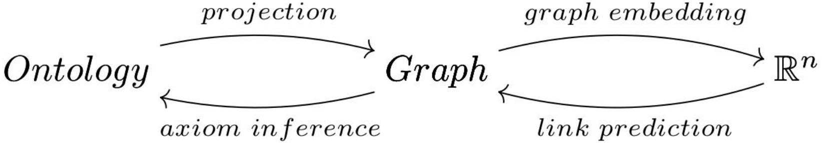

Some graph projections are not embeddings:

- injectivity: $C \sqsubseteq \exists R.D$ and $C \sqsubseteq \forall R.D$
- totality: $\neg C \sqsubseteq \neg D$

Properties of graph projections

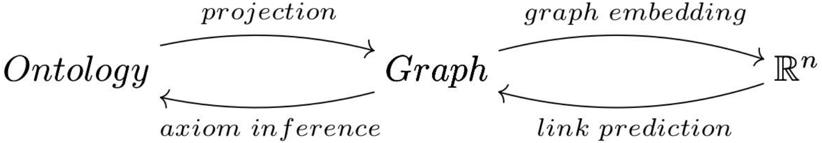

Some graph projections are not embeddings:

- injectivity: $C \sqsubseteq \exists R.D$ and $C \sqsubseteq \forall R.D$
- totality: $\neg C \sqsubseteq \neg D$
- not invertible $\rightarrow$ no direct axiom inference

Evaluation: neural deductive inference $C \sqsubseteq D$

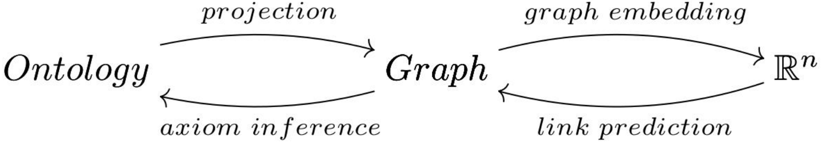

Can graph-based ontology embeddings compute entailments?

Evaluation: neural deductive inference $C \sqsubseteq D$

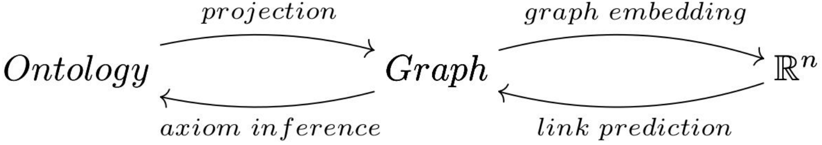

Can graph-based ontology embeddings compute entailments?

- Gene Ontology, deductive closure
- using TransE ($h + r \approx t$) for link prediction

Evaluation: neural deductive inference $C \sqsubseteq D$

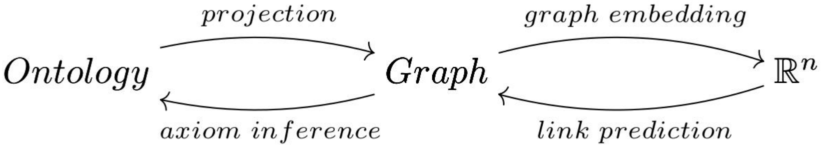

Can graph-based ontology embeddings compute entailments?

- Gene Ontology, deductive closure
- using TransE ($h + r \approx t$) for link prediction

Method	*GO*									
	MR		H@1		H@10		H@100		AUC	
	R	F	R	F	R	F	R	F	R	F
Onto2Graph	**2952.63**	**2947.81**	0.41	0.41	0.85	0.85	12.22	12.32	**94.22**	**94.23**
OWL2Vec*	4519.32	4514.51	**1.70**	**1.98**	**9.45**	**10.03**	**33.40**	**33.87**	91.15	91.16
RDF	4163.59	4158.77	0.36	0.36	2.82	2.85	9.98	10.03	91.85	91.86

Graph projections (again)

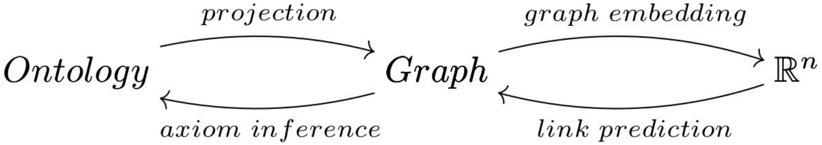

Graph projections (again)

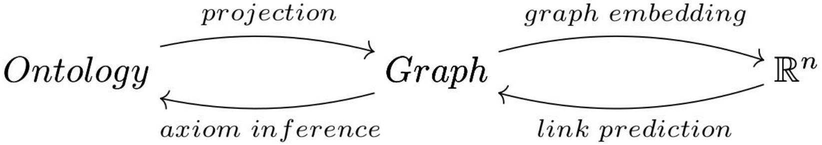

- Lattice structure
- LatE — Lattice-preserving graph embeddings

CatE — Lattice-preserving embeddings

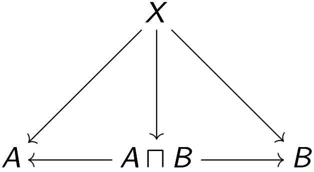

(a) Intersection

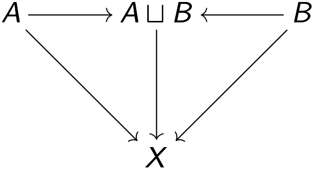

(b) Union

Figure: Lattice representations of complex concept descriptions.

Negation

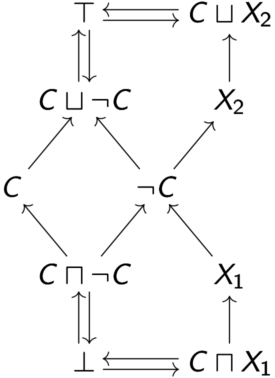
$\top$
$C \sqcup X_2$
$C \sqcup \neg C$
X_2
C
$\neg C$
$C \sqcap \neg C$
X_1
$\bot$
$C \sqcap X_1$

CatE — Results

Table: Prediction of axioms $C(a)$ in ORE1 dataset

Method	MR	MRR	Hits@3	Hits@10	Hits@100	AUC
ELEmbeddings	105	0.12	0.08	0.22	0.87	0.99
Box^{2}EL	122	0.10	0.08	0.18	0.70	0.98
FALCON	603	0.02	0.00	0.02	0.34	0.92
CatE	**37**	**0.18**	**0.10**	**0.51**	**0.96**	**0.99**

Entailment and approximate entailment

- semantic entailment: $T \models \phi$ iff $Mod(T) \subseteq Mod(\{\phi\})$
  - $Mod(T) = \{\mathcal{A} | \text{for all } \psi \in T\text{: } \mathcal{A} \models \psi\}$

Entailment and approximate entailment

- semantic entailment: $T \models \phi$ iff $Mod(T) \subseteq Mod(\{\phi\})$
  - $Mod(T) = \{\mathcal{A} | \text{for all } \psi \in T\text{: } \mathcal{A} \models \psi\}$
- approximate semantic entailment:
  - $T \models_f \phi$ iff $Mod_f(T) \subseteq Mod(\{\phi\})$ with $Mod_f(T) \subsetneq Mod(T)$
  - for ~~all~~ $\psi \in T$: $\mathcal{A} \models \psi$ (approximate models)
  - for all $\psi \in T$: $\mathcal{A} \models^{\alpha} \psi$ for some $\alpha \geq \tau$ (fuzzy models)
- generating approximate models of T through algorithm P:
  - P is *faithful* if it converges to a model of T
  - P is *representation complete* if it can converge to any model of T

Model-generating embeddings

- a detour through conceptual spaces...
- semantics of OWL relies on interpretation (Σ-algebras)
- an interpretation is a *model* of T if all $C \sqsubseteq D \in T$ are satisfied

Name	Syntax	Semantics
top	$\top$	$\Delta^{\mathcal{I}}$
bottom	$\bot$	$\emptyset$
nominal	$\{a\}$	$\{a^{\mathcal{I}}\}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
existential restriction	$\exists r.C$	$\{x \in \Delta^{\mathcal{I}} \mid \exists y \in \Delta^{\mathcal{I}} : (x, y) \in r^{\mathcal{I}} \land y \in C^{\mathcal{I}}\}$
generalized concept inclusion	$C \sqsubseteq D$	$C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$
role inclusion	$r_1 \circ ... \circ r_n \sqsubseteq r$	$r_1^{\mathcal{I}} \circ ... \circ r_n^{\mathcal{I}} \subseteq r^{\mathcal{I}}$

EL Embeddings

- given a knowledge base T with signature $\Sigma(T)$
- aim: find embedding $f_e : \Sigma(T) \mapsto \mathbb{R}^n$ s.t. $f_e(\Sigma(T))$ is a model of T ($f_e(\Sigma(T)) \models T$)
    - map symbols into $\mathbb{R}^n$ $==$ construct a model in $\mathbb{R}^n$
- guaranteed existence:
    - any consistent $\mathcal{EL}^{++}$theory has infinite models
    - any consistent $\mathcal{EL}^{++}$theory has models in $\mathbb{R}^n$ (theorem of Löwenheim-Skolem, upwards)

Algorithm: loss functions

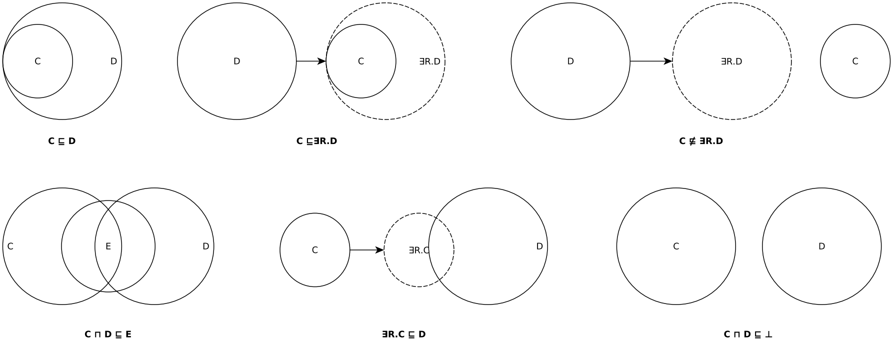

EL Embeddings

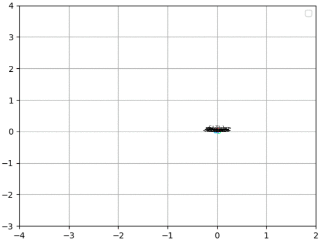

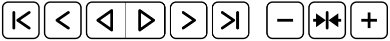

Theorem (Faithfulness)

Let T be a theory in $\mathcal{EL}^{++}$. If $\gamma \leq 0$ and $loss_n(\eta(T)) = 0$ then T has a model.

Kulmanov et al., 2019

$Male$	$\sqsubseteq Person$
$Female$	$\sqsubseteq Person$
$Father$	$\sqsubseteq Male$
$Mother$	$\sqsubseteq Female$
$Father$	$\sqsubseteq Parent$
$Mother$	$\sqsubseteq Parent$
$Female \sqcap Male$	$\sqsubseteq \bot$
$Female \sqcap Parent$	$\sqsubseteq Mother$
$Male \sqcap Parent$	$\sqsubseteq Father$
$\exists hasChild.Person$	$\sqsubseteq Parent$
$Parent$	$\sqsubseteq Person$
$Parent$	$\sqsubseteq \exists hasChild.\top$

Limitations of EL Embeddings

- intersection of spheres is no sphere (closure w.r.t. intersection)
  - axis-aligned boxes instead of circles

Limitations of EL Embeddings

- intersection of spheres is no sphere (closure w.r.t. intersection)
  - axis-aligned boxes instead of circles
- approximate entailment ($T \models \phi$ iff $Mod_f(T) \subseteq Mod(\{\phi\})$))
  - generate multiple models, aggregation function over models

Limitations of EL Embeddings

- intersection of spheres is no sphere (closure w.r.t. intersection)
  - axis-aligned boxes instead of circles
- approximate entailment ($T \models \phi$ iff $Mod_f(T) \subseteq Mod(\{\phi\})$))
  - generate multiple models, aggregation function over models
- not every model can be represented, no representation completeness:
  - $X \equiv A \sqcap B$, $Y \equiv B \sqcap C$, $Z \equiv A \sqcap C$
    $X \sqcap Y \sqsubseteq \bot$, $Y \sqcap Z \sqsubseteq \bot$, $X \sqcap Z \sqsubseteq \bot$

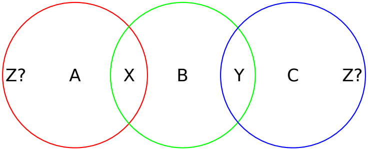

Limitations of EL Embeddings

- intersection of spheres is no sphere (closure w.r.t. intersection)
  - axis-aligned boxes instead of circles
- approximate entailment ($T \models \phi$ iff $Mod_f(T) \subseteq Mod(\{\phi\})$))
  - generate multiple models, aggregation function over models
- not every model can be represented, no representation completeness:
  - $X \equiv A \sqcap B$, $Y \equiv B \sqcap C$, $Z \equiv A \sqcap C$
    $X \sqcap Y \sqsubseteq \bot$, $Y \sqcap Z \sqsubseteq \bot$, $X \sqcap Z \sqsubseteq \bot$

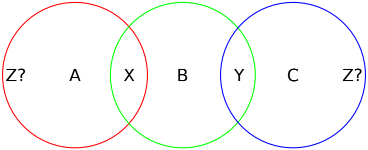

- may be a fundamental limitation of geometric embeddings
  - cones, non-Euclidean geometry, fractals?

FALCON

- fuzzy logic with differentiable t-norm θ and negator ν
- generating approximate models using two functions:
  - embedding f_e and interpretation $f_{\mathcal{I}}$
- recursive forward to parse axioms

FALCON

- fuzzy logic with differentiable t-norm θ and negator ν
- generating approximate models using two functions:
  - embedding f_e and interpretation $f_{\mathcal{I}}$
- recursive forward to parse axioms

$$m(x, C^{\mathcal{I}}) = \sigma(MLP(f_e(C), f_e(x)))$$

$$m((x, y), R^{\mathcal{I}}) = \sigma(MLP(f_e(x) + f_e(R), f_e(y)))$$

FALCON

- fuzzy logic with differentiable t-norm θ and negator ν
- generating approximate models using two functions:
  - embedding f_e and interpretation $f_{\mathcal{I}}$
- recursive forward to parse axioms

$$m(x, C^{\mathcal{I}}) = \sigma(MLP(f_e(C), f_e(x)))$$

$$m((x, y), R^{\mathcal{I}}) = \sigma(MLP(f_e(x) + f_e(R), f_e(y)))$$

$$m(x, (\neg C)^{\mathcal{I}}) = \nu(m(x, C^{\mathcal{I}}))$$

FALCON

- fuzzy logic with differentiable t-norm θ and negator ν
- generating approximate models using two functions:
  - embedding f_e and interpretation $f_{\mathcal{I}}$
- recursive forward to parse axioms

$$m(x, C^{\mathcal{I}}) = \sigma(MLP(f_e(C), f_e(x)))$$
$$m((x, y), R^{\mathcal{I}}) = \sigma(MLP(f_e(x) + f_e(R), f_e(y)))$$

$$m(x, (\neg C)^{\mathcal{I}}) = \nu(m(x, C^{\mathcal{I}}))$$

$$m(x, (C_1 \sqcap C_2)^{\mathcal{I}}) = \theta(m(x, C_1^{\mathcal{I}}), m(x, C_2^{\mathcal{I}})$$

FALCON

- fuzzy logic with differentiable t-norm θ and negator ν
- generating approximate models using two functions:
  - embedding f_e and interpretation $f_{\mathcal{I}}$
- recursive forward to parse axioms

$$m(x, C^{\mathcal{I}}) = \sigma(MLP(f_e(C), f_e(x)))$$

$$m((x, y), R^{\mathcal{I}}) = \sigma(MLP(f_e(x) + f_e(R), f_e(y)))$$

$$m(x, (\neg C)^{\mathcal{I}}) = \nu(m(x, C^{\mathcal{I}}))$$

$$m(x, (C_1 \sqcap C_2)^{\mathcal{I}}) = \theta(m(x, C_1^{\mathcal{I}}), m(x, C_2^{\mathcal{I}})$$

$$m(x, (\exists R.D)^{\mathcal{I}}) = \max_{y \in \Delta} \theta(m(y, D^{\mathcal{I}}), m((x, y), R^{\mathcal{I}}))$$

FALCON

- approximate semantic entailment: $T \models_f \phi$ iff $Mod_f(T) \subseteq Mod(\phi)$
    - generate multiple approximate models
    - can sample $\mathbb{R}^n$ for anonymous individuals $\Rightarrow$ allows for infinite domains
- Theorem 1 (faithfulness): If $loss = 0$, FALCON finds model of T.
- Theorem 2 (representation completeness): If $\mathcal{I}$ is a finite model of T, FALCON can "represent" it with $loss = 0$.

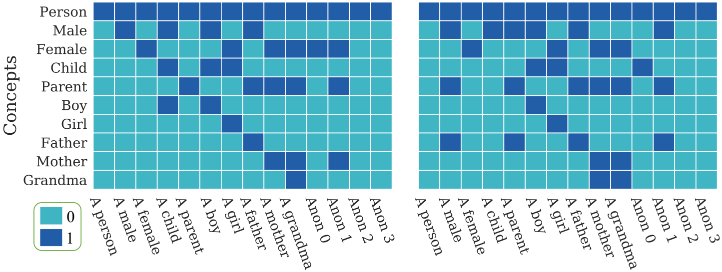

mOWL

- Python library for machine learning with Semantic Web (OWL) ontologies
- ontology embeddings, zero-shot, constrained optimization
- contains
  - graph generation
  - graph embedding
  - approximate models
  - semantic entailment
- Algorithms written in Python + Scala, tuned for performance

`https://github.com/bio-ontology-research-group/mowl`

Knowledge-based methods in bioinformatics

- protein function prediction
- diagnosing rare diseases (gene–disease associations)
- interactions between drugs/proteins and proteins

DeepGOZero: putting it all together

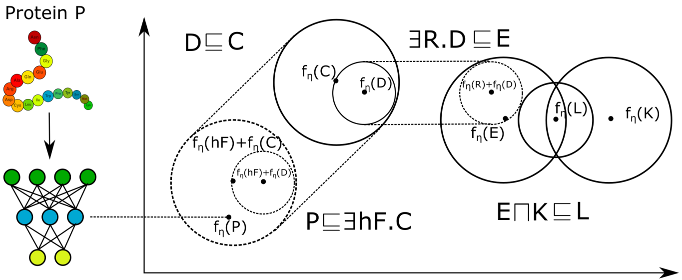

- knowledge-based protein function prediction system
- webserver with over 100 million jobs completed since 2017
- multi-class, multi-label, neuro-symbolic zero-shot prediction

DeepGOZero: putting it all together

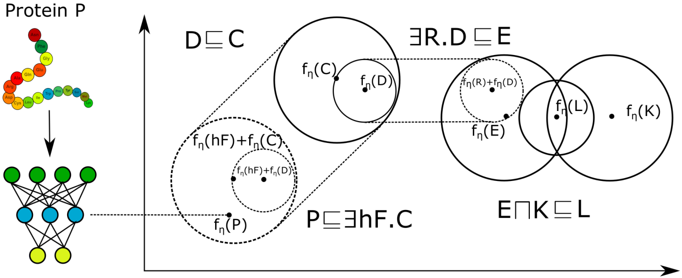

- EL Embeddings to construct *one* model $\mathcal{A}$ of GO
- MLP projects protein embedding p into EL Embedding space
  - supervised: $p \sqsubseteq \exists \text{hasFunction}.GO$ true in $\mathcal{A}$
- inference is a test for (approximate) "truth":
  $y'_c = \sigma(f_\eta(p) \cdot (f_\eta(hF) + f_\eta(c))^T + r_\eta(c))$

Zero shot protein function prediction

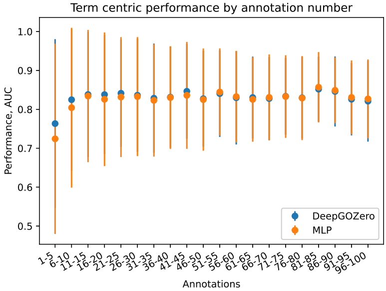

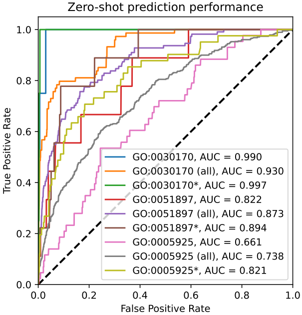

- GO class compositions are preserved in the embedding space!

DeepGO Semantic Entailment

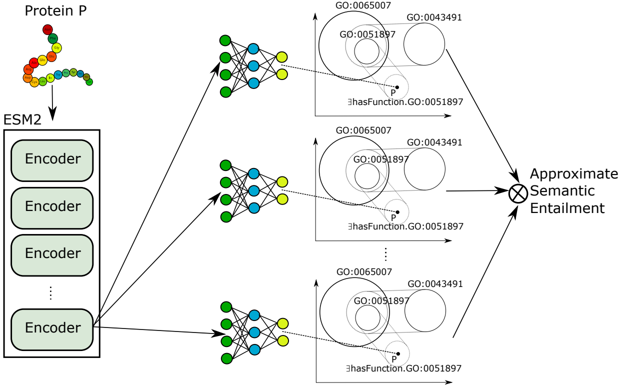

- Truth in single model → statements may be true by chance
- $T \models \phi$ iff $Mod(T) \subseteq Mod(\phi)$

$$y'_c = SE_{i=1}^{N}(\sigma(f_\eta^i(p) \cdot (f_\eta^i(hF) + f_\eta^i(c))^T + r_\eta(c)^i)$$

Kulmanov et al., Nature Machine Intelligence, 2024

DeepGO Semantic Entailment

Table: Prediction results for Molecular Function on the UniProtKB/Swiss-Prot dataset

Method	Fmax	Smin	AUPR	AUC
Naive	0.321	14.568	0.180	0.500
MLP	0.321	14.606	0.195	0.500
MLP (ESM2)	0.517	12.197	0.508	0.830
DeepGOCNN	0.404	13.741	0.365	0.749
DeepGOZero	0.483	12.722	0.444	0.749
DeepGraphGO	0.416	14.077	0.357	0.673
DeepGO-SE	**0.554**	11.681	**0.552**	**0.874**
DeepGOGAT-SE	0.525	**11.137**	0.523	0.861

DeepGO Semantic Entailment

Table: Prediction results for Biological Process on the UniProtKB/Swiss-Prot dataset

Method	Fmax	Smin	AUPR	AUC
Naive	0.294	43.934	0.195	0.500
MLP	0.295	43.914	0.210	0.499
MLP (ESM2)	0.423	39.721	0.388	0.864
DeepGOCNN	0.334	42.912	0.275	0.686
DeepGOZero	0.343	42.857	0.284	0.643
DeepGraphGO	0.354	42.100	0.303	0.736
DeepGO-SE	0.432	39.419	0.401	0.864
DeepGOGAT-SE	0.435	39.123	0.404	**0.876**
DeepGOGATMF-SE	**0.448**	**37.299**	**0.428**	0.831
DeepGOGATMF-SE-Pred	0.444	39.098	0.409	0.855

Summary

Feigenbaum, 1977

[The domain-specific knowledge] plays a critical role in organizing and constraining search. The theme is that in the knowledge is the power.

- biology is a knowledge-based discipline
  - supplies use-cases to develop knowledge-based methods in AI
  - immediate impact, generalizable to other domains
- neuro-symbolic methods: allow neural networks to generate and reason over world models
  - focus on “knowledge” and “reasoning”
- formal properties: soundness / faithfulness, completeness / representation completeness, convergence

Acknowledgements

Thank you

Amyloid beta
Protein classified with blood
coagulation.

A Semantic Haiku

generated from the UniProt Knowledgebase

`http://borg.kaust.edu.sa`
robert.hoehndorf@kaust.edu.sa