Explaining Protein-Protein Interaction Predictions with Genetic Programming

Rita T. Sousa, Sara Silva, Catia Pesquita
LASIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal

Motivation

Explainability is crucial to support the adoption of machine learning as a tool for scientific discovery. In the biomedical domain, ontologies and knowledge graphs (KG) are a unique opportunity to explore domain knowledge, but most KG-based approaches employ graph embeddings, which are not explainable. Since similarity assessment is a natural explanatory mechanism, an alternative explanatory strategy is to use KGs to measure the semantic similarity (SS) between entities in the graph.

We address the explainability problem for protein-protein interaction (PPI) prediction by using genetic programming (GP) algorithms over KG-based SS values to generate global and interpretable explanations.
Since SS can be computed using different portions of the KG to reflect different semantic aspects, we propose that SS for different semantic aspects can provide more granular explanations with higher information content.

Methods

PPI prediction is cast as a binary classification task that takes as input the Gene Ontology (GO) KG and a PPI dataset containing a set of protein pairs that interact.

Ontology and KG-based SS afford a unique opportunity to explore domain knowledge and support explainability

Computing KG-based SS for each Semantic Aspect

SS scores were computed for the 50 subgraphs rooted in the direct subclasses of GO roots (removing aspects with potential bias for PPI).

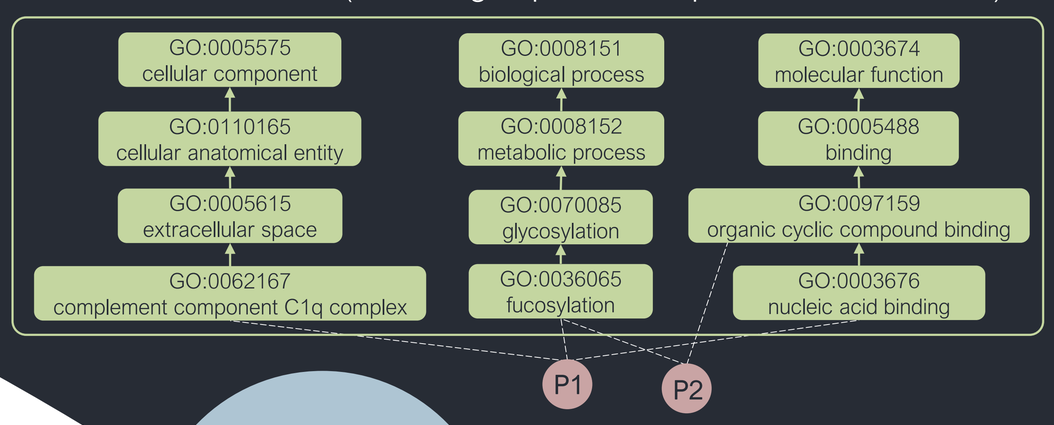

GP can be fine-tuned for interpretability

Evolving a GP Model

GP	GP6x
No depth penalization. 6 basic operators.	Penalize solutions with a depth > 6. Interpretable operators in the biological context (max,min,+,-).

Results

We performed a 10-fold cross-validation:

	GP	GP6x
Median weighted average of F-measures	0.875	0.866
Median number of nodes	49	17

With small differences, all GP6x models consider maximum similarities of multiple semantic aspects with a majority describing biological processes. This corroborates prior knowledge that for two proteins to interact they usually participate in the same biological processes.

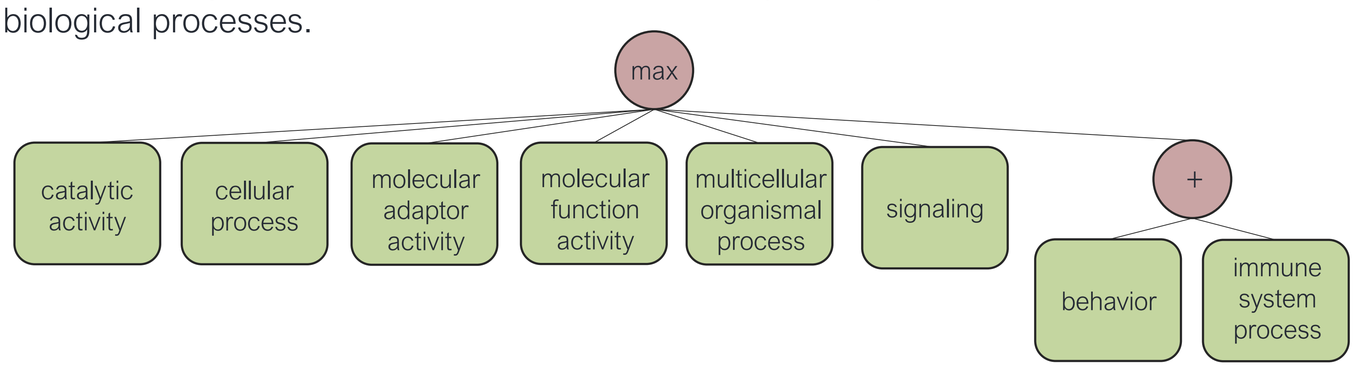

Four protein pairs were chosen among well-predicted positive pairs (+/+), well-predicted negative pairs (-/-), wrongly-predicted positive pairs (+/-), and wrongly-predicted negative pairs (-/+) and the input SS values were analysed:

40S ribosomal protein S12 and 40S ribosomal protein S10 (+/+)

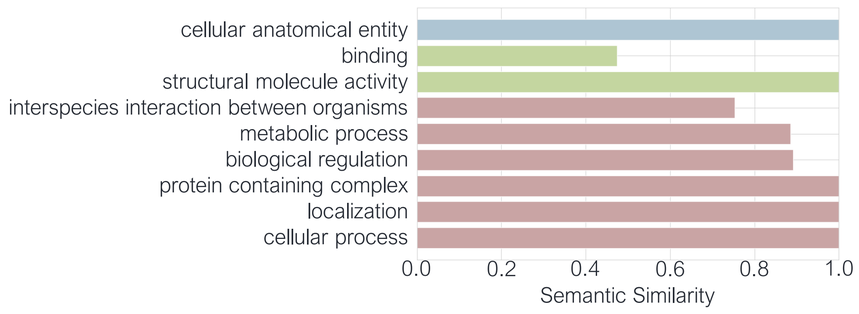

Kinetochore-associated protein 1 and Tubulin β-6 chain (-/-)

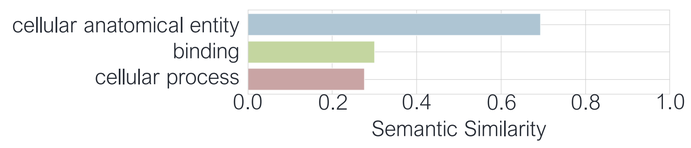

S100-A10 protein and neuroblast differentiation-associated protein AHNAK (+/-)

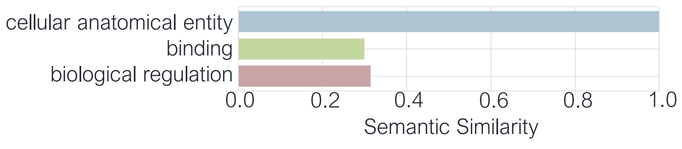

Both proteins participate in the development of the intracellular membrane. The misclassification can be justified by the incomplete annotation of these proteins.

Protransforming growth factor α and Disks large homolog 2 (-/+)

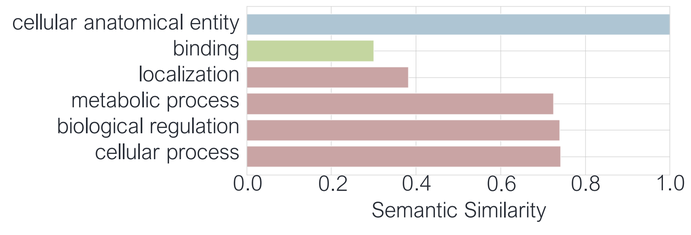

The literature describes interactions between proteins of the same family of the pair, indicating that this is likely a true but still unknown interaction.

Prediction explanations can uncover new knowledge

Conclusions

GP can produce global and interpretable explanations. The performance of the more interpretable GP methods is lower, but what little they sacrifice in performance is more than gained in explainability.

Explainability can be key to uncover issues with the underlying data and even pose new hypothesis.

Acknowledgements: This work was funded by FCT through LASIGE Research Unit (UIDB/00408/2020, UIDP/00408/2020); projects GADgET (DSAIPA/DS/0022/2018) and BINDER (PTDC/CCI-INF/29168/2017); PhD grant SFRH/BD/145377/2019.