Project Assignments

SML Lab, University of Trento

Advanced Topics in Machine Learning and Optimization — 2024-25

Predicting SAT hardness via GNNs

Assignment

- **SAT**isfiability is the problem of checking whether a logic formula is satisfiable by a truth assignment. Graph Neural Networks (GNNs) have been applied to many SAT-related task, from end-to-end SAT solving to feature extraction.
- **Goal**: use **GNNs** for predicting the runtime of SAT instances.
- **Applications**: solver selection in portfolio approaches, theoretical understanding of SAT hardness
- The student is asked to:
  - Create a dataset of formulas/runtime; see (Liu et al., 2024).
  - Train (and evaluate) the GNNs for predicting the runtime.
  - *(option)* Analyzing "similar" formulas exhibiting significantly different runtimes.
  - *(option)* Train on multiple SAT solvers.

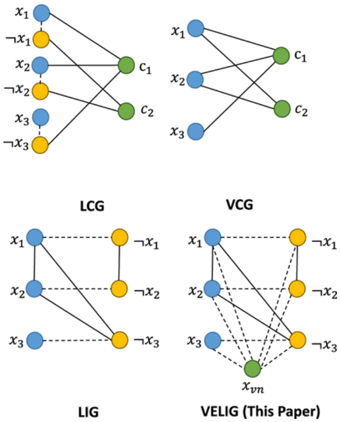

Multiple graph representation of a formula.

Notes

- **Contact**: Antonio Longa, Stefano Teso, Paolo Morettin
- Extensible to thesis!

Approximating volume computations via GNNs

Assignment

- Beyond SAT, Graph Neural Networks (GNNs) have been used for approximating other computationally hard problems. We turn to the problem of computing the volume of a **convex polytopes** $\Delta = \bigwedge_i \mathbf{A}_i \mathbf{x} \leq b$
- **Goal**: find a suitable graphical representation of Δ, then train a GNN for approximating $vol(\Delta)$.
- **Applications**: constrained probabilistic inference
- The student is asked to:
  - Create a dataset of $\langle \Delta_k, vol(\Delta_k) \rangle$ pairs
  - Implement a procedure for encoding polytopes as graphs. Train (and evaluate) the GNNs.
  - *(option)* Experiment with different encodings.
  - *(option)* Compare with sequence-based models (provided).

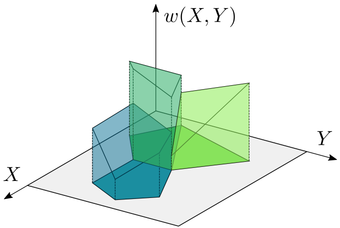

Polytopes encoding a complex probability distribution.

Notes

- **Contact**: Veronica Lachi Paolo Morettin
- Extensible to thesis!

Counterfactuals with Conformal Predictors

Assignment

- *Conformal prediction* is a framework for uncertainty quantification which can be used to build decision support systems that improve experts' accuracy in classification tasks with theoretical guarantees via *prediction sets* (Straitouri et al., 2023). However, the connection between conformal prediction and explanations is still unexplored.
- The student is asked to:
  - Generate a synthetic dataset for multi-label classification and train a classifier.
  - Implement a *conformal predictor* to extract *prediction sets*.
  - Implement a *counterfactual generating* technique optimizing *sparsity* or *proximity*.
  - Analyze the relationship between *sparsity* and *proximity* and the variations in the *prediction sets*.

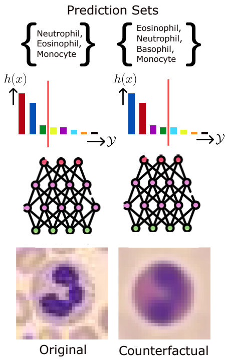

Notes

- **Contact**: Cesare Barbera, Giovanni De Toni
- Extensible to thesis!

Explanatory Interaction: Ask Smarter Questions!

Assignment

- Explanatory Interactive Learning (Schramowski et al., 2020) is great for *deconfounding* models, however it builds on very simple active learning strategies for selecting items to be labelled. **Can we do better?**
- The student is asked to:
  - Take existing explanatory debugging code and integrate a state-of-the-art algorithm for confounding-aware query selection.
  - Evaluate on confounded data whether this strategy brings the intended benefits.

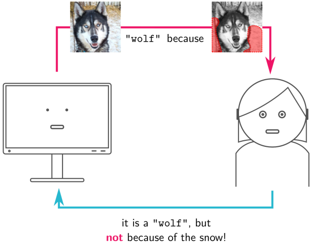

Explanatory interactive learning (XIL)

Notes

- **Contact**: Stefano Teso
- Extensible to thesis!

Learning to Guide Users with CBMs, RLHF-style

Assignment

- In **Learning to Guide** (LTG), a model is learned to produce *textual guidance* useful for guiding human decision makers. This requires fine-tuning an LLM, and in turn a lot of expensive supervision. **Can we do better with CBMs**?
- The student is asked to:
  - Read up on learning to guide (Banerjee et al., 2024).
  - Design a (simple) classification task
  - Design and train surrogate model that scores the CBM's explanations based on how much they are helpful for a down-stream decision maker.
  - Teach the CBM to output more useful explanations, **RLHF style**!

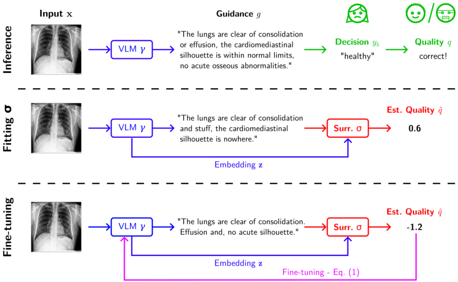

Notes

- **Contact**: Debodeep Banerjee, Burcu Sayin, Stefano Teso
- Extensible to thesis!

Improving Concepts with Vision-Language Models

Assignment

- **Concept-based Models** (CBMs) learn to map input images to high-level concepts, such as "`fluffy dog`" and "`red car`". These concepts may however not be high quality.
- Can we exploit **Vision-Language Models** (VLMs) to obtain feedback about properties that learned concepts ought to satisfy? (Srivastava et al., 2024)
- The student is asked to:
  - Design a (simple) classification task in which concepts learned by a stock CBM is not great.
  - Query a VLM about properties of learned concepts ("should the ball be red?") and align the CBM accordingly.

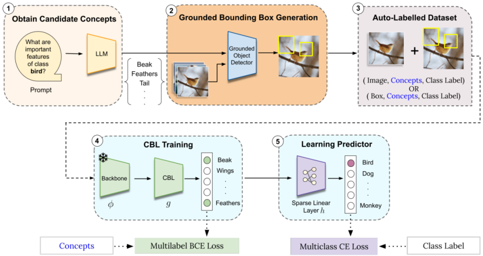

Notes

- **Contact**: Emanuele Marconato, Stefano Teso
- Extensible to thesis!

Argumentative interaction with LLMs in Legal Sector

Assignment

- Unchecked use of LLMs in law is risky; involving humans in the loop of decision-making in legal tasks is more ethical and practical.
- The student is asked to:
  - Develop an LLM-powered Intelligent Assistant (IA) to augment human decision-makers in the legal field. The IA will not make a final decision but rather help the legal expert make a final decision through an argumentative interaction.

Notes

- **Contact**: Burcu Sayin, Andrea Passerini
- Extensible to thesis!

Shortcuts in NeSy Models with Structured Knowledge

Assignment

- Inductive Logic Programming (Evans and Grefenstette, 2018) (ILP) allows to learn logical rules from examples and background knowledge.
- NeSy models (that learns both concepts and knowledge) are prone to reasoning shortcuts and algorithmic shortcuts. What happens when knowledge is more complex?
- The student is asked to:
  - Design an end-to-end NeSy architecture that uses differentiable ILP;
  - Create a logical task that admits reasoning and algorithmic shortcuts;
  - Compare this model with existing ones. Is it more robust to shortcuts?

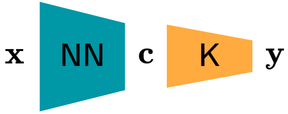

x: 

c: [A, B]

K:

A	B	Y
0	0	0
0	1	1
1	0	1
1	1	0

Notes

- **Contact**: Samuele Bortolotti, Emanuele Marconato
- Extensible to thesis!

Can Neuro-Symbolic models cheat at IQ tests?

Assignment

- NeSy architectures combine learning and perception with reasoning. Sometimes, they solve the task by learning concepts with unexpected or unclear semantics. These are **reasoning shortcuts**.
- Can NeSy models do so in IQ tests like **Raven matrices**?
- The student is asked to:
  - Read up on reasoning shortcuts (Bortolotti et al., 2024) and Raven matrices (Zhang et al., 2019).
  - Design a (simplified) Raven data set for one of the NeSy models available in (Bortolotti et al., 2024).
  - Learn a model on your data and evaluate whether it learns any reasoning shortcuts.

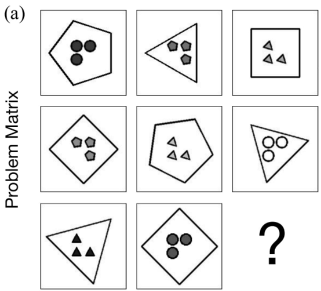

A Raven matrix. What should the missing square be?

Notes

- **Contact**: Samuele Bortolotti, Emanuele Marconato, Stefano Teso
- Extensible to thesis!

Algorithmic Reasoning in Large Language Models

Assignment

- Surprisingly, Large Language Models perform really well on various tasks they are not trained for.
- Arithmetic Reasoning (e.g., sum, product, ...) is known to be a serious limitation for these models (Dziri et al., 2023).
- We want to test the hypothesis that LLMs do not learn **the intended reasoning algorithm**: the cheapest variant is discovered, which does not generalize.
- The student is asked to:
  - Design tasks of arithmetic reasoning.
  - Use chain-of-thought (COT) techniques to obtain the "reasoning" algorithm of LLMs (Dziri et al., 2023)
  - Validate through **Mechanistic Intertpretability** whether the COT is faithful (Wu et al., 2023). .

```
function multiply (x[1..p], y[1..q]):
  // multiply x for each y[i]
  for i = q to 1
    carry = 0
    for j = p to 1
      t = x[j] * y[i]
      t += carry
      carry = t // 10
      digits[j] = t mod 10
    summands[i] = digits

  // add partial results (computation not shown)
  product = Σ_{i=1}^{q} summands[q+1-i] · 10^{i-1}
  return product
```

$A(\mathbf{x})$

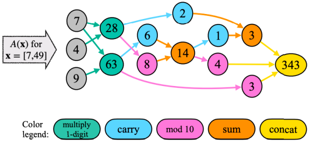

Notes

- **Contact**: Emanuele Marconato, Samuele Bortolotti
- Extensible to thesis!

GNN node classification for homophilic and heterophilic labels

Assignment

- The scientific literature contains somewhat inconclusive and potentially conflicting statements regarding the role of *Label Homophily* (the property that two connected nodes are likely to have the same label) for the success of GNN node classifiers.
- The student is asked to:
  - Review recent research literature and critically evaluate existing experimental evaluations of classification performance under different levels of homophily
  - Design new and improved experimental evaluations, e.g. by creating customized synthetic data
  - Explore possible combinations of GNN classifiers with other techniques to optimize performance both under Homo- and Heterophily

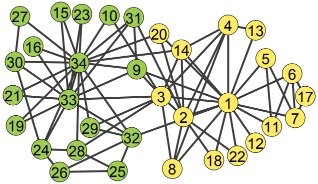

Notes

- **Contact**: Andrea Passerini Manfred Jaeger,
- Extensible to thesis!

Graph Neural Networks for Relational Databases

Assignment

- Databases are ubiquitous, yet deep learning often underperforms compared to tree-based models (e.g., boosting and random forests) in database-related tasks. A recent trend involves converting relational databases into graphs and applying Graph Neural Network models to them. However, these graphs exhibit distinctive features, such as being multipartite, high-diameter, and highly heterogeneous.
- The student is asked to:
  - Test standard GNN models on relational databases benchmark. (RelBench[a])
  - Investigate potential limitations of GNN models when applied to multipartite graphs.
  - The project can be tailored based on the student's preferences, allowing for either (a) a more theoretical focus, for those interested in conceptual exploration. (b) An empirical focus, for students who prefer practical experimentation and analysis.
- Further details can be discussed in a brainstorming session, either in person or online, with Francesco, Antonio, and Veronica.

[a] https://relbench.stanford.edu/

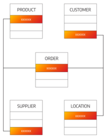

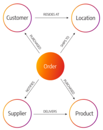

Notes

- **Contact**: Francesco Ferrini, Antonio Longa, Veronica Lachi
- Extensible to thesis!

Assignment

- Although Message Passing is the most common framework for learning relational data, it has notable limitations, including Oversmoothing (Chen et al., 2020), Oversquashing (Di Giovanni et al., 2023), Underreaching (Errica et al., 2023), and challenges with Robustness (Günnemann, 2022).
- It remains unclear which of these issues most significantly impacts the training dynamics and generalization performance of GNNs.
- The student is tasked with:
  - Investigating how (some of) these phenomena evolve during training and whether they correlate with generalization error.
  - Examining how (some of) these phenomena evolve during training and whether they impact the quality of the explanations (Yuan et al., 2022; De Luca et al., 2024).

Notes

- **Contact**: Vincenzo Marco De Luca
- Extensible to thesis!

Graph Neural Networks: Can they Generalize with limited data?

Assignment

- Numerous regularization techniques have been proposed for Graph Neural Networks, with most focusing on topological-level adjustments (e.g., rewiring, dropping (Fang et al., 2023).
- While these strategies have shown positive effects on reducing generalization error, the GNN literature still lacks a thorough comparison with prominent Few-Shot Learning (Satorras and Estrach, 2018) strategies.
- The student is tasked with:
  - Benchmark various topological regularizers and Few-Shot Learning strategies on the most famous graph datasets to evaluate their potential.

Notes

- **Contact**: Vincenzo Marco De Luca
- Extensible to thesis!

Exam modality

Project work

- Select one of the projects from the previous slides (or discuss with the teacher for custom projects)
- Complete it and prepare a report summarizing the methodology used and the results obtained
- After completing the assignment send it via email to the (first) contact person for the project
- Subject: ADVML2024
- Attachment: name_surname.zip containing:
    - the report (named report.pdf)
    - the code you wrote
    - the requirements needed to run the code

NOTE

- No group work
- Preliminary versions of the report can be sent for feedback
- The project is discussed asynchronously as soon as it is completed

References

Banerjee, D., Teso, S., Sayin, B., and Passerini, A. (2024). Learning to guide human decision makers with vision-language models. *arXiv preprint arXiv:2403.16501*.

Bortolotti, S., Marconato, E., Carraro, T., Morettin, P., van Krieken, E., Vergari, A., Teso, S., and Passerini, A. (2024). A neuro-symbolic benchmark suite for concept quality and reasoning shortcuts. In *The Thirty-Eighth Annual Conference on Neural Information Processing Systems*, pages 1–44. ACM.

Chen, D., Lin, Y., Li, W., Li, P., Zhou, J., and Sun, X. (2020). Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 3438–3445.

De Luca, V. M., Longa, A., Passerini, A., and Liò, P. (2024). xai-drop: Don't use what you cannot explain. *arXiv preprint arXiv:2407.20067*.

Di Giovanni, F., Giusti, L., Barbero, F., Luise, G., Lio, P., and Bronstein, M. M. (2023). On over-squashing in message passing neural networks: The impact of width, depth, and topology. In *International Conference on Machine Learning*, pages 7865–7885. PMLR.

Dziri, N., Lu, X., Sclar, M., Li, X. L., Jiang, L., Lin, B. Y., West, P., Bhagavatula, C., Bras, R. L., Hwang, J. D., Sanyal, S., Welleck, S., Ren, X., Ettinger, A., Harchaoui, Z., and Choi, Y. (2023). Faith and fate: Limits of transformers on compositionality.

Errica, F., Christiansen, H., Zaverkin, V., Maruyama, T., Niepert, M., and Alesiani, F. (2023). Adaptive message passing: A general framework to mitigate oversmoothing, oversquashing, and underreaching. *arXiv preprint arXiv:2312.16560*.

Evans, R. and Grefenstette, E. (2018). Learning explanatory rules from noisy data.

Fang, T., Xiao, Z., Wang, C., Xu, J., Yang, X., and Yang, Y. (2023). Dropmessage: Unifying random dropping for graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 4267–4275.

Günnemann, S. (2022). Graph neural networks: Adversarial robustness. *Graph neural networks: foundations, frontiers, and applications*, pages 149–176.

Liu, J., Xiao, W., Cheng, H., and Shi, C. (2024). Graph neural network based time estimator for sat solver. *International Journal of Machine Learning and Cybernetics*, pages 1–12.

Satorras, V. G. and Estrach, J. B. (2018). Few-shot learning with graph neural networks. In *International conference on learning representations*.

Schramowski, P., Stammer, W., Teso, S., Brugger, A., Herbert, F., Shao, X., Luigs, H.-G., Mahlein, A.-K., and Kersting, K. (2020). Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nature Machine Intelligence*, 2(8):476–486.

Srivastava, D., Yan, G., and Weng, T.-W. (2024). Vlg-cbm: Training concept bottleneck models with vision-language guidance. *arXiv preprint arXiv:2408.01432*.

Straitouri, E., Wang, L., Okati, N., and Rodriguez, M. G. (2023). Improving expert predictions with conformal prediction. In *International Conference on Machine Learning*, pages 32633–32653. PMLR.

Wu, Z., Geiger, A., Icard, T., Potts, C., and Goodman, N. (2023). Interpretability at scale: Identifying causal mechanisms in alpaca. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S., editors, *Advances in Neural Information Processing Systems*, volume 36, pages 78205–78226. Curran Associates, Inc.

Yuan, H., Yu, H., Gui, S., and Ji, S. (2022). Explainability in graph neural networks: A taxonomic survey. *IEEE transactions on pattern analysis and machine intelligence*, 45(5):5782–5799.

Zhang, C., Gao, F., Jia, B., Zhu, Y., and Zhu, S.-C. (2019). Raven: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5317–5327.