Project Assignments

Andrea Passerini
andrea.passerini@unitn.it

Advanced Topics in Machine Learning and Optimization

Improve separation of local GNN explanations

Description

- GLGExplainer is a Global GNN Explainer. It learns a set of concepts using two regularization losses and the prediction of the GNN to explain. In some circumstances, these losses are not enough to separate local explanations.
- The goal is to implement a technique (Autoencoders, self-supervised learning, or unsupervised pertaining) to better promote the separation of local explanations.
- We mainly refer to the MNIST dataset. The goal is to classify even and odd numbers represented as graphs

Notes

- **Contact**: Steve Azzolin, Antonio Longa
- **Material**: Paper
- **Extensible to thesis**

NeurASP vs DeepProbLog

Description

- **NeurASP:** A neuro-symbolic framework that extends Answer Set Programming (ASP) with neural networks
- Compare NeurASP vs DeepProbLog on Sudoku solving
- Retrain the NeurASP model on Sudoku and collect accuracy (as done in the paper)
- Train a DeepProbLog model on the same data and collect accuracy
- Compare accuracy, show predictions, and report training time of both frameworks

Notes

- **Contact**: Gianluca Apriceno
- **Extensible to thesis**
- **References**: paper, repository

Robustness of Algorithmic Recourse

Description

- Current frameworks providing algorithmic recourse assume that users can carry out the suggested interventions precisely. In reality, users execute the interventions in a *noisy* manner.
- In this project, the aim is to investigate the robustness of learned *recourse policies* under *noisy* actuation.
- Given the *FARE* framework, evaluate its robustness to noisy interventions and, optionally, devise about a new robust learning objective.

Notes

- **Contact**: Giovanni De Toni
- **Resources**: Paper
- **Extensible to thesis**

Alignment metric for concept representations

Description

- Choose a deep latent variables model and two disentanglement datasets from the literature
- Train the model on different scenarios of generative factor supervision (e.g. increasing supervision, noisy examples, weak supervision, etc.)
- **Define** a metric for **alignment**, which combines an existing **disentanglement** measure + **your measure** of **monotonicity.**
- Compare results from the two datasets and assess the pros and cons of the **alignment metric**

Notes

- **Contact**: Emanuele Marconato
- **Extensible to thesis**

Identifiability of Representations through NeSy

Description

- Create a supervised task where **disentangled** ground-truth generative factors define a NeuroSymbolic task
- Train DeepProblog on a bunch of different scenarios (i.e. few *vs* many classes to predict, absence *vs* presence of generative factors supervision)
- Evaluate **disentanglement** performances and assess the identifiability of the encoder
- **Q:** When does DeepProbLog achieve **unsupervised** concept discovery?

Notes

- **Contact**: Emanuele Marconato
- **Extensible to thesis**

Implement Weakly-supervised GlanceNets

Description

- **GlanceNets** are concept-based models that combine high-performance with sound interpretability
- Implement a loss enabling them to learn from weak supervision, e.g., pairs of images that differ only in one concept
- Compare them to the original in terms of performance and interpretability
- Optionally, implement an extra monotonicity term geared toward better interpretability, and measure its impact

Notes

- **Contact**: Stefano Teso, Emanuele Marconato
- **Materials**: paper and code
- **Extensible to thesis**

Active Learning for Neuro-Symbolic Integration

Description

- **Semantic Probabilistic Layers** (SPLs) ensure neural network predictions are always consistent with given symbolic knowledge.
- Implement an Active Learning loop of an SPL-augmented neural network for hierarchical or structured prediction.
- Evaluate the performance of a few Active Learning selection strategies with and without the SPL layer.
- Optionally, implement and evaluate an advanced batch-based selection strategy.

Notes

- **Contact**: Stefano Teso
- **Collaboration**: Antonio Vergari, University of Edinburgh.
- **Material**: paper and code.
- **Extensible to thesis**

Neuro-Symbolic Machine Unlearning

Description

- Consider an online learning setting in which the model predictions are constrained by **symbolic knowldge that changes over time**. How to teach the model to behave consistently when this happens?
- Implement one or more losses for forgetting obsolete knowledge on top of a neural network for, e.g., hierarchical image classification.
- Evaluate that it works as intended.

Notes

- **Contact**: Stefano Teso, Andrea Bontempelli
- **Material**: paper
- **Extensible to thesis**

DeepProbLog vs Deep Network

Assignment

- Consider the MNIST multi-digit addition example we saw in the lecture
- Design a purely neural architecture that predicts the result of the addition (e.g. CNN+LSTM, but feel free to invent)
- Compare results of the neural architecture with those of DeepProbLog in the same setting:
  - supervision only on the result of the addition
  - generalization to longer numbers
- Try training the neural network with a larger training set than the one used for DeepProbLog (check how many examples are need to match DeepProbLog performance)

Notes

- **Contact**: Andrea Passerini
- **Can be selected multiple times**

DeepProbLog Tasks

Description

- Start from the MNIST multi-digit addition example we saw in the lecture
- Modify it to address *octal* division between two numbers (assume second is an integer divisor of first)
- Always assume (as in the example) that training instances have one-digit numbers, test instances multiple digit numbers

Notes

- **Contact**: Andrea Passerini

Graph forecasting

Description

- Similar to time series forecasting: predict future development of graph evolution process
- Special case: predict development of information cascade in social network
- Train recurrent graph neural network models on information cascade data (e.g. Twitter)
- Apply trained network to predict future development of observed initial cascade

Notes

- **Contact**: Andrea Passerini
- **Collaboration**: Manfred Jaeger, Aalborg University, DK
- Possibility for internship abroad
- **Extensible to thesis**

Debugging Interpretable NNs for Medical Decision Making Support

Description

- X-Ray scans of the lungs can identify patterns that indicate problems due to infections (e.g. by covid-19)
- Deep architectures for medical image classification often tend to learn spurious correlations that do not generalize (confounders)
- The task is that of training a confounder-free interpretable NN model (e.g. a PPNet) by:
  - Applying ProtoPDebug to lung image classification
  - Setting up a user study where users provide feedback to prevent ProtoPDebug from learning confounders

Notes

- **Contact**: Andrea Passerini, Andrea Bontempelli
- **Extensible to thesis**

Neural Production Systems

Description

- Production Systems are inference tools developed in symbolic AI and based on **rules**
- There is a growing interest in the Deep Learning community in exploring rule-based generalization
- In this project, the goal is to develop a Production System using only a deep network
- We will start from existing literature on this topic and generalize the current methods introducing an unsupervised learning of the system's rules

Notes

- **Contact**: Andrea Passerini
- **Collaboration**: Enver Sangineto, University of Modena and Reggio Emilia
- **Extensible to thesis**

Assignment

- Select one of the projects from the previous slides (or discussed 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 andrea.passerini@unitn.it
- Subject: ADVML2022
- 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