Some Applications of ML to Adaptive Network Security

Adaptive/Stream Learning Models for NetSec

- Adaptive learning algorithms **trained on labelled data, using ADWIN**

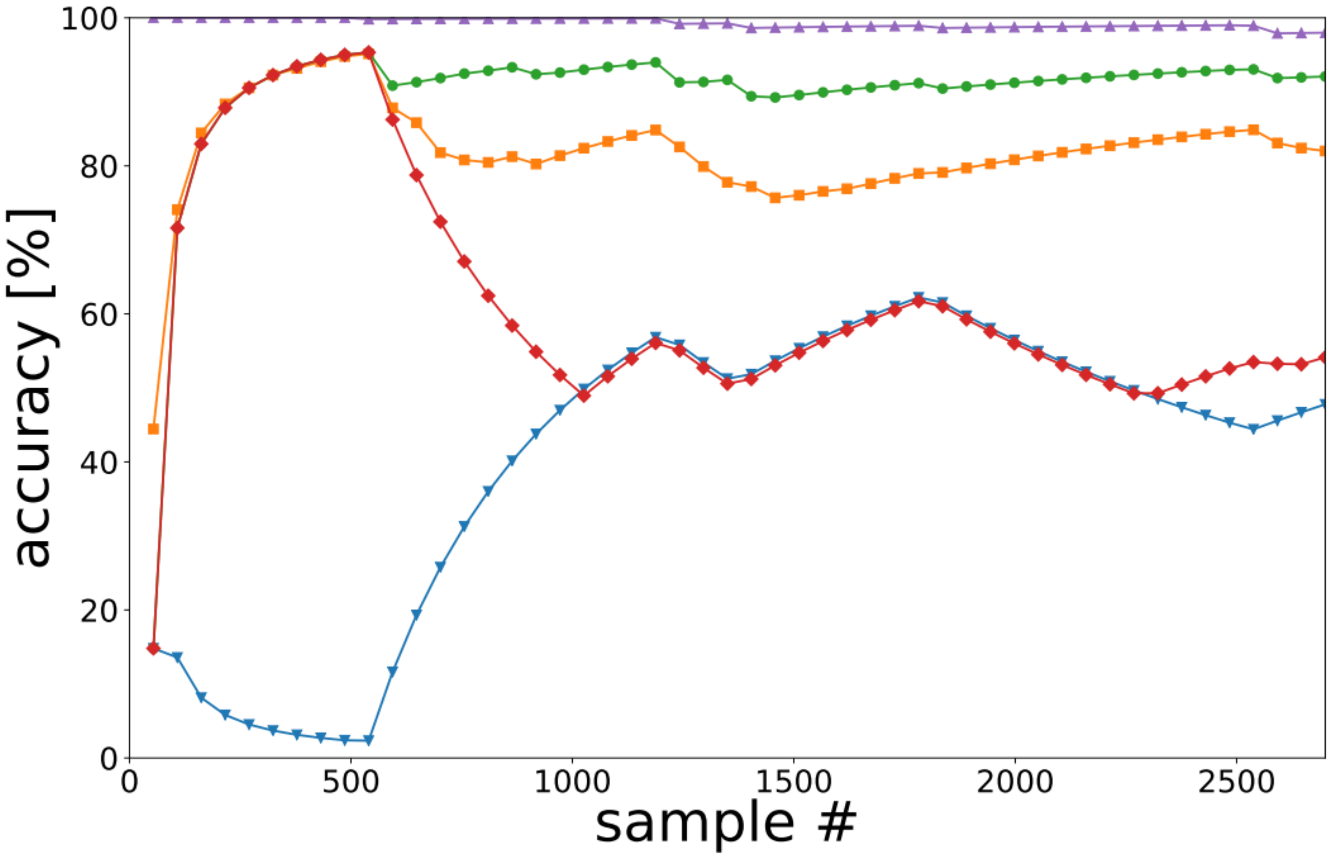

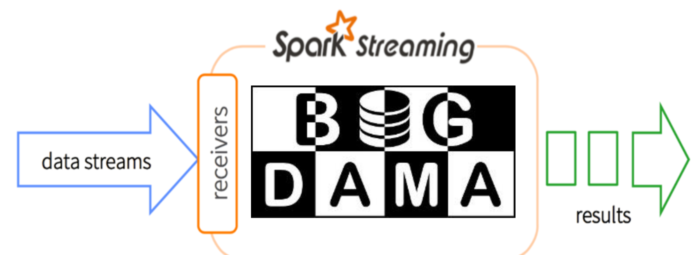

Stream-based Learning Models Performance

- Multiple stream machine learning models, using ADWIN
- Detection accuracy, ***normalized to batch-based algorithms*** performance

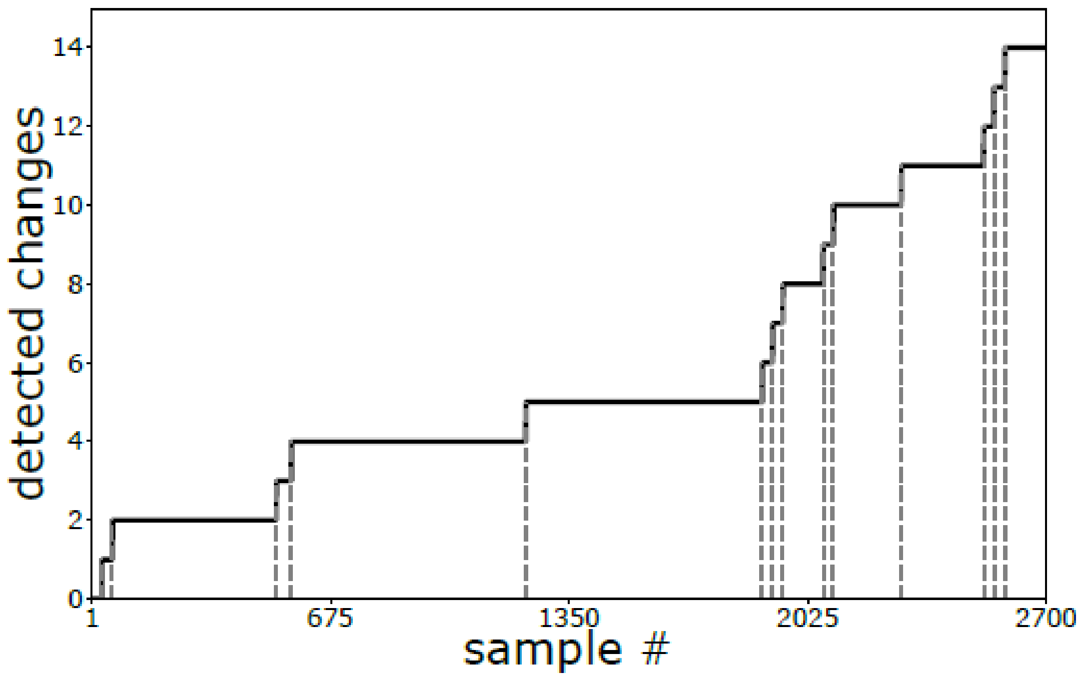

Figure 1: ***Page-Hinkley Concept Drift Detection.*** Changes in the dataset distribution detected by the Page-Hinkley test. Detected changes are marked with dashed lines.

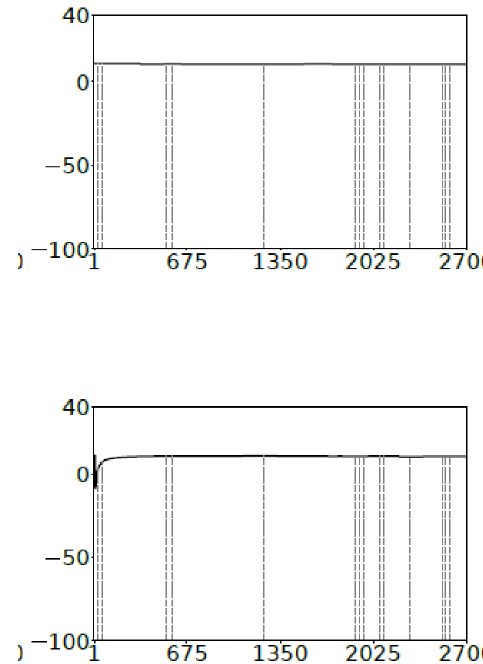

(d) *SGD*

Stream-based Learning Models Performance

- Multiple stream machine learning models, using ***fixed windowing***
- AUC (ROC curve), ***normalized to batch-based algorithms*** performance
- **Different window sizes tested**

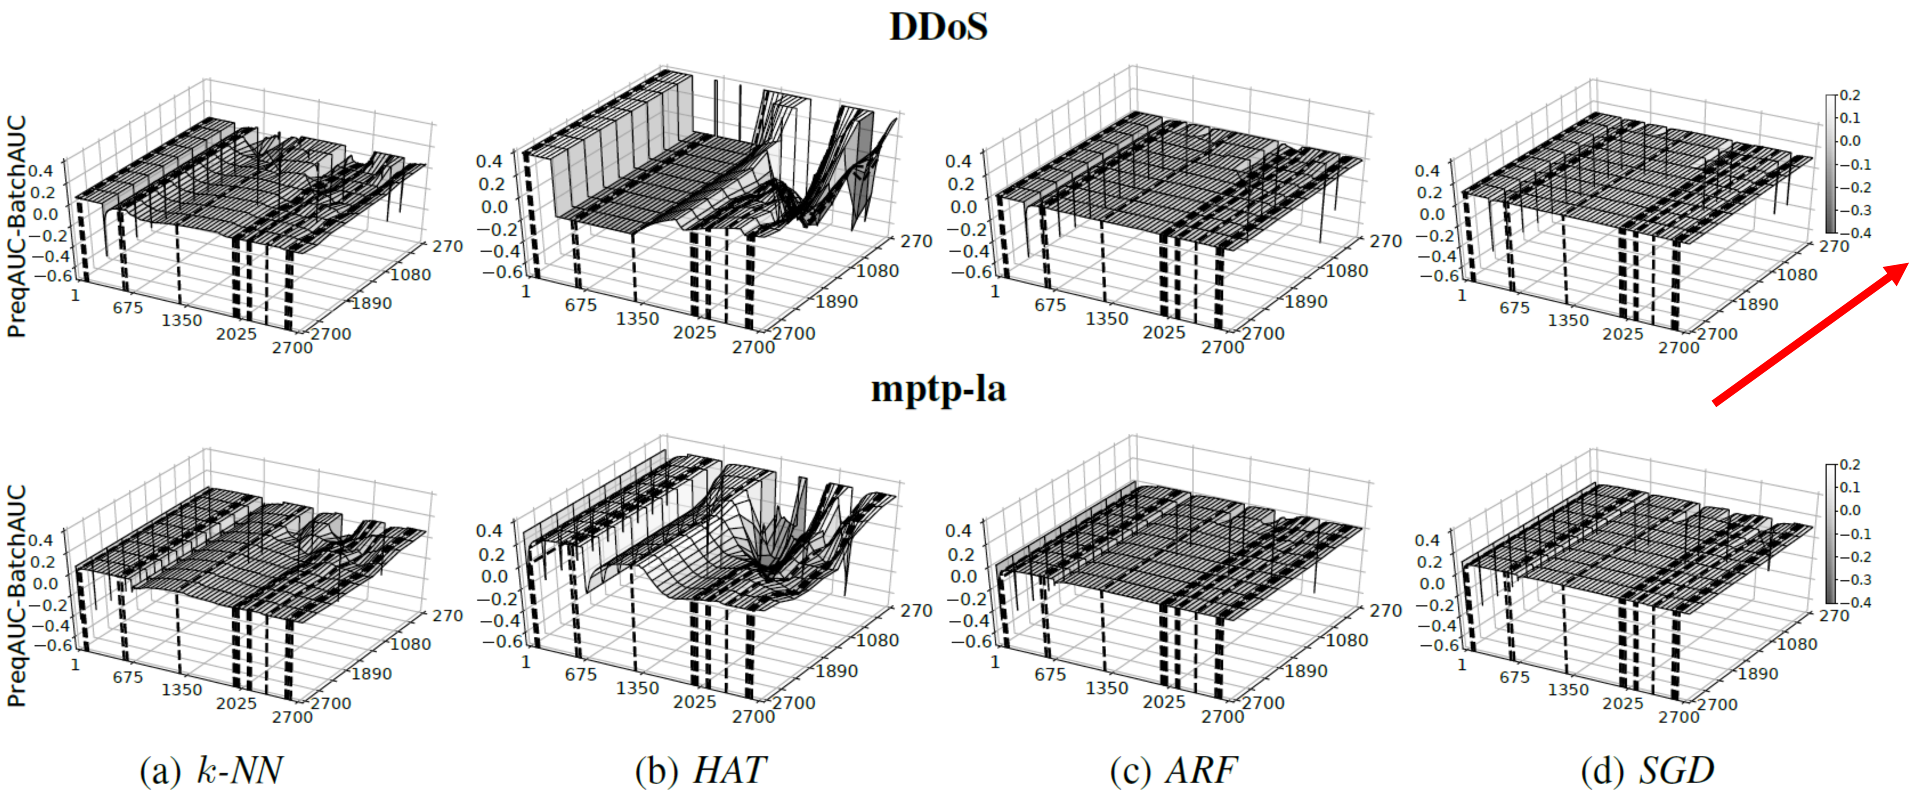

Improving Stream-based Active Learning by Reinforcement (RAL)

- How do we deal with the **limited amount of labeled data?**
- **Active Learning (AL):** aims at labelling only the most informative samples
- **AL** can be applied to the **streaming scenario**, to complement previous approaches and **reduce the amount of labeled data**

- RAL – improves stream-based AL by Reinforcement Learning (RL)
  - AL bases its decisions based EXCLUSUVELY on model uncertainty
  - **RAL permits to additionally learn in a feedback loop**, based on the **effectiveness of the requested labels**
  - **Reward** in case asking oracle was informative (models would have predicted wrong label)
  - **Penalty** otherwise

RAL Principles and Components

- RAL is based on an ensemble of models
- RAL makes use of **contextual-bandit algorithms** (EXP4) to tune the decision powers of the different models depending on their behavior
- RAL uses a **ε-greedy approach** to handle concept drift and **improve the exploration/exploitation trade-off**

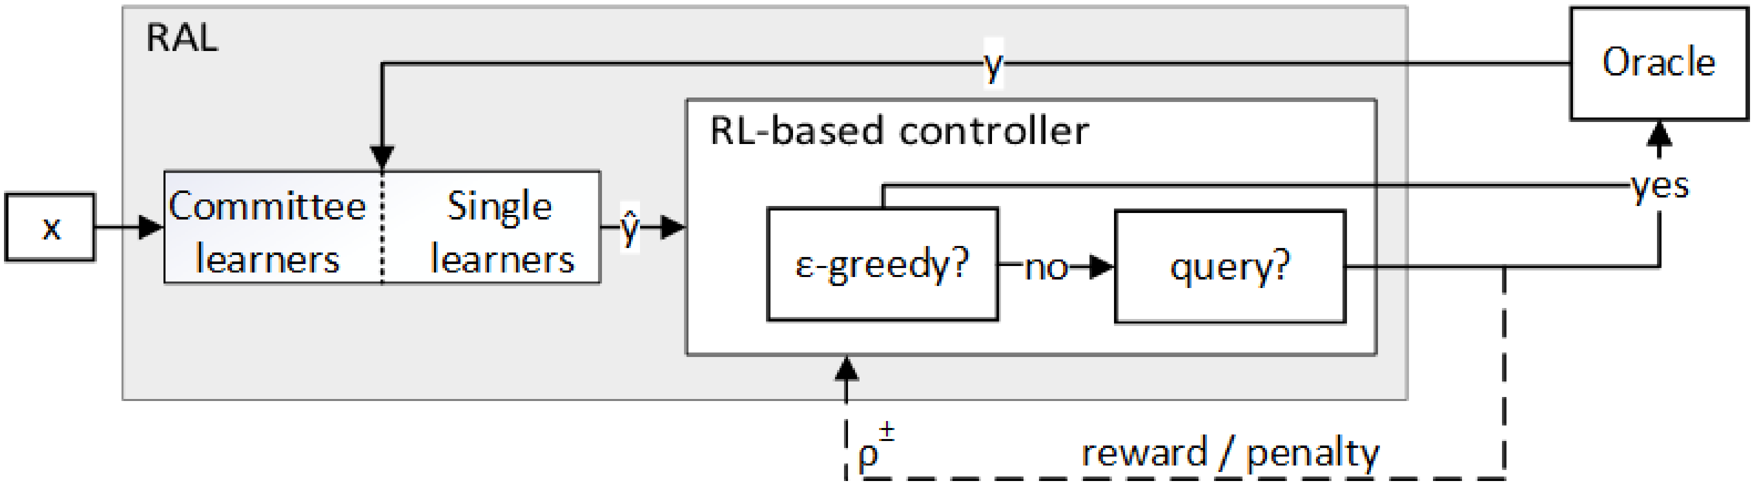

RAL Principles and Components

- The **querying decision** (ask or not for a label) is taken **based on model prediction uncertainty** and a **threshold**
- Each algorithm in the ensemble (committee) gives its advice, based on its prediction uncertainty
- RAL takes into account the decisions of the members + their decision power
- Obtained **feedback influences the querying threshold**:
  - In case of **penalty**, the threshold decreases…..otherwise, it **slightly increases**

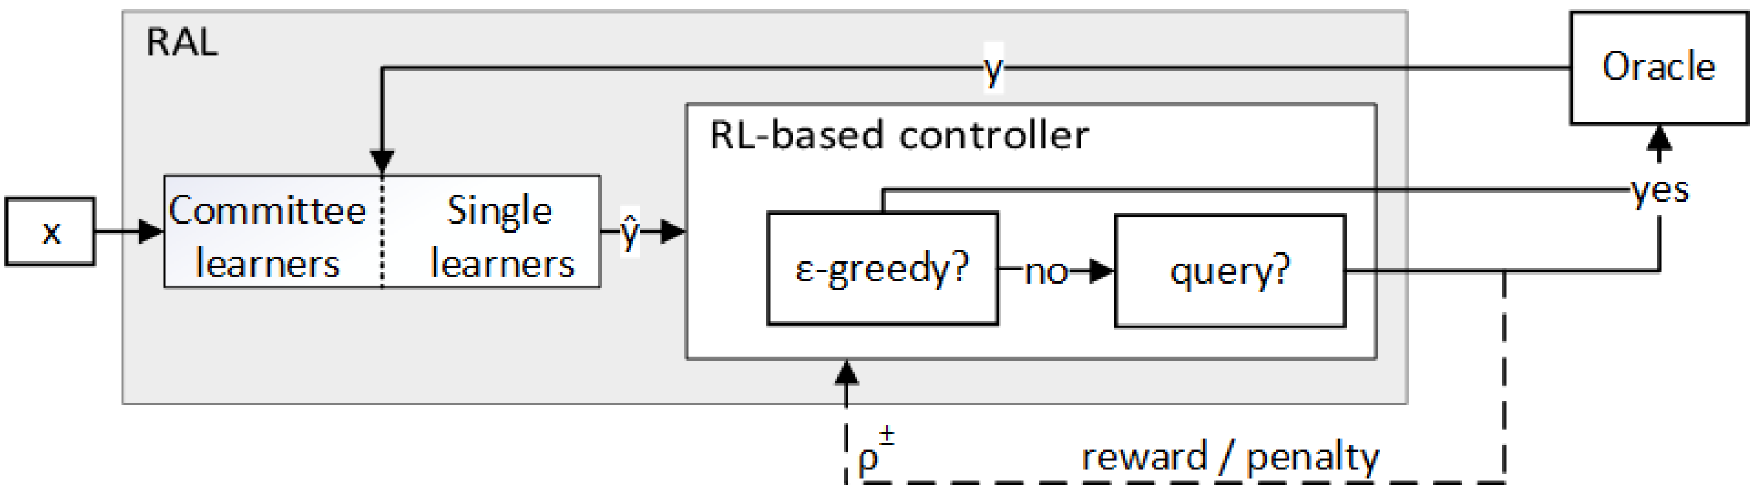

RAL Evaluation vs. State of the Art

- **RAL vs** RVU (**Randomized Variable Uncertainty**) and simple **random sampling** (RS)

- Evaluation on data extracted from **MAWILab – in the wild network security**

- We divide each dataset into three consecutive parts:
  - **Initial training set** (variable size)
  - **Validation set** (last 30%), to evaluate the classifiers
  - **Streaming set** (remaining part of the dataset), for picking samples to learn from

RAL Evaluation vs. State of the Art – Prediction Accuracy

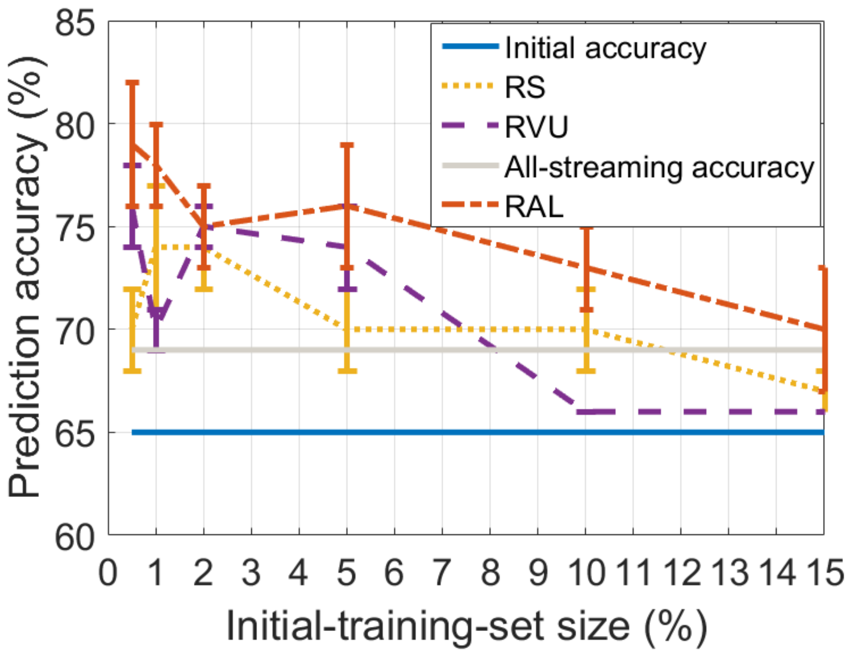

Flood attack

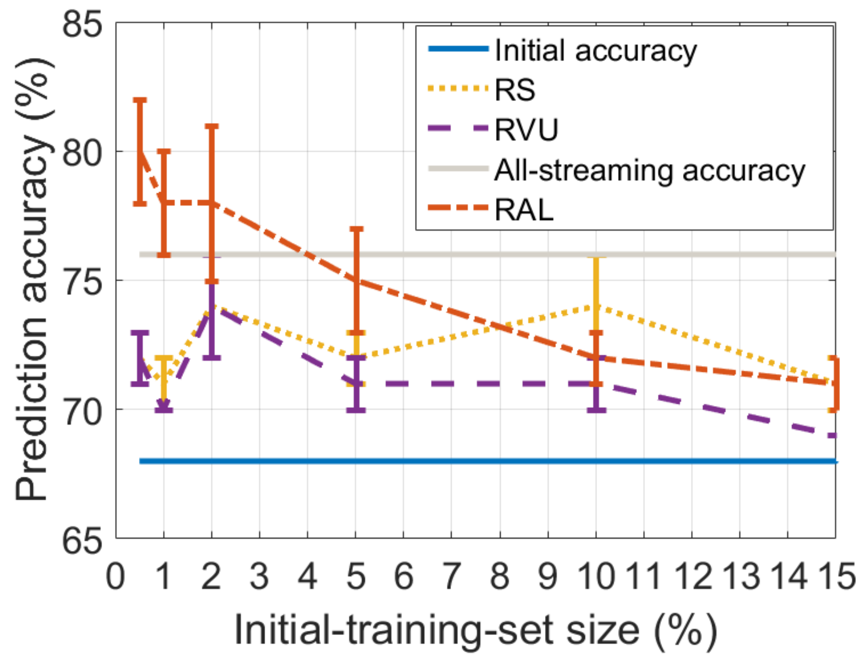

Netscan attack

RAL Evaluation vs. State of the Art – Querying Cost

RAL

RVU

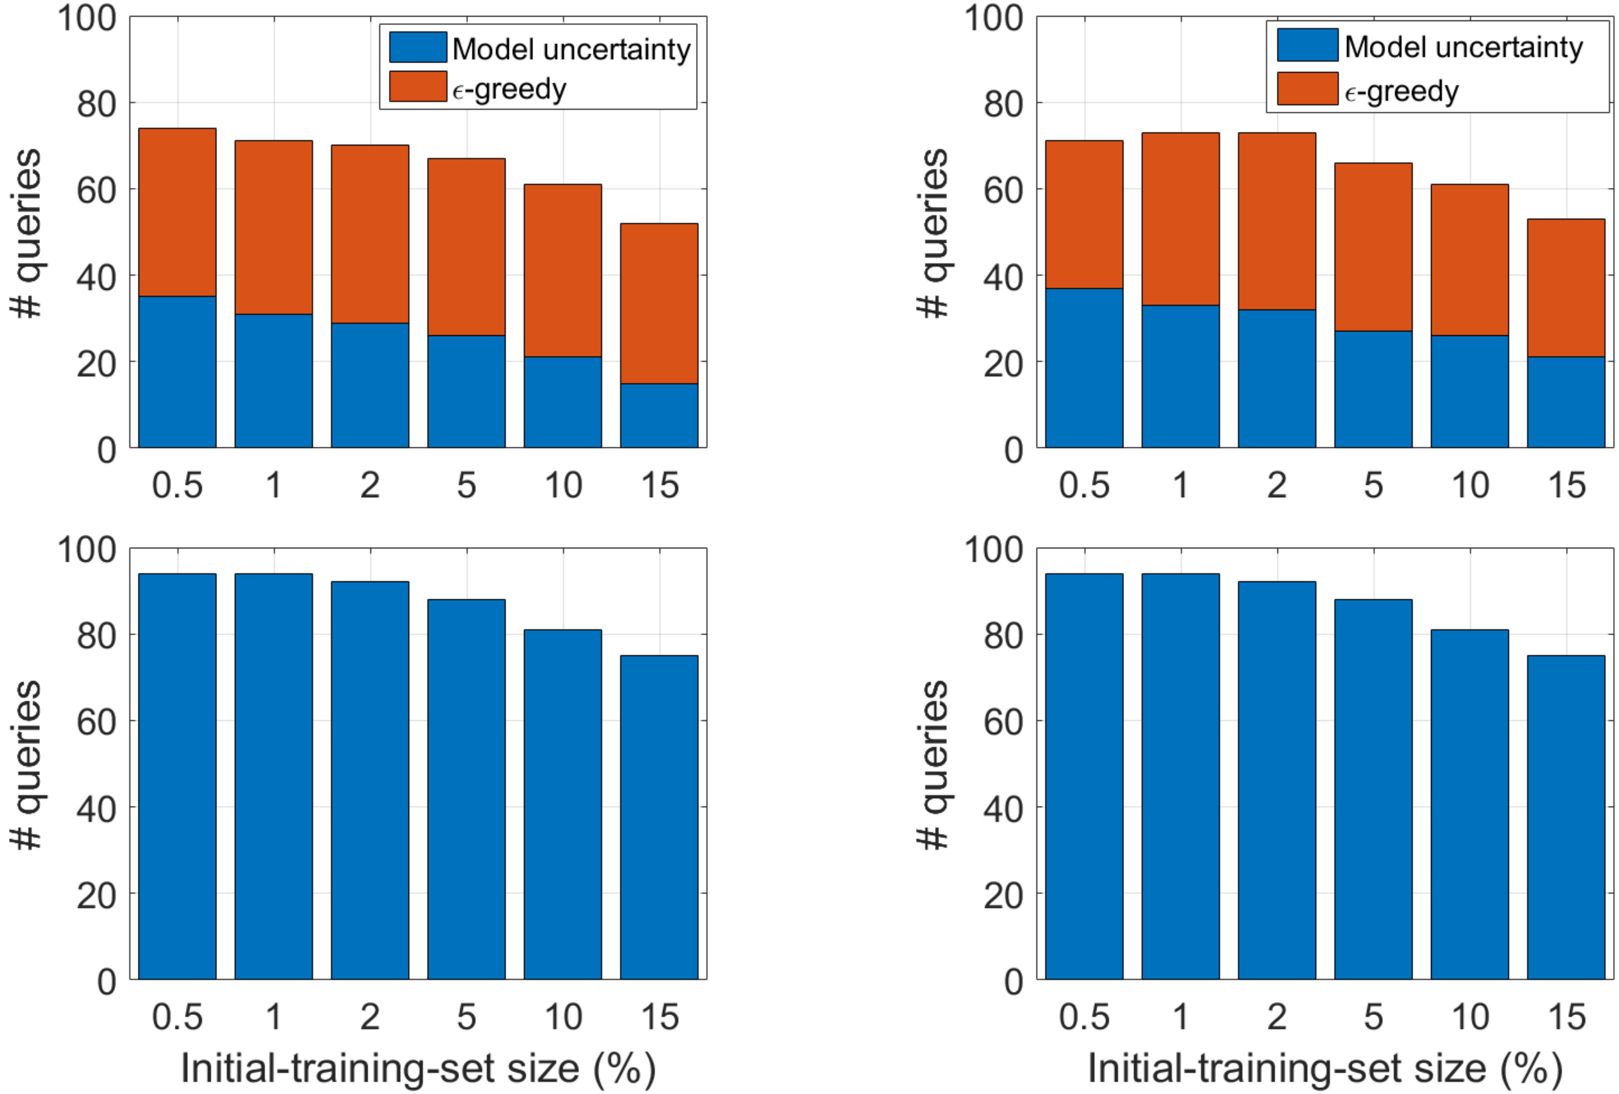

Flood attack

Netscan attack

So What’s Next?

- We’re still far from ***making AI immediately applicable***
  - Limitations of learning process, data, models
  - Lack of generalization
  - Continual learning challenges – catastrophic forgetting and transfer
  - Lack of real knowledge generation – building simple mappings is *easy*
  - Portability of models to real deployments – *plug & play?*
- ***Effective Machine Learning*** – a mix of interesting challenges:
  - Transfer learning
  - Explainable AI (XAI)
  - Multi-task learning
  - Meta learning
  - Hierarchical learning
- And back right to the start: **the successful application of AI to network measurement problems is still on an early stage**

https://bigdama.ait.ac.at/

http://mobiqoe.ait.ac.at/

Thanks

Dr. Pedro Casas

Data Science @Digital Insight Lab

AIT Austrian Institute of Technology @Vienna

pedro.casas@ait.ac.at

http://pcasas.info