# Real-time Semantic Computer Vision for Co-robotics

## Towards Realistic Predictors

**Pei Wang, Nuno Vasconcelos**

Statistical Visual Computing Lab, UC San Diego

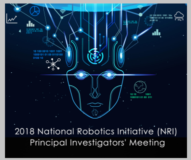


SVCL UCSD

# Motivation

- computer vision algorithms are optimistic
- they attempt to recognize/ detect/segment/ all instances
- without regard to how hard the task is
- this is not what humans do!

cristiano ronaldo

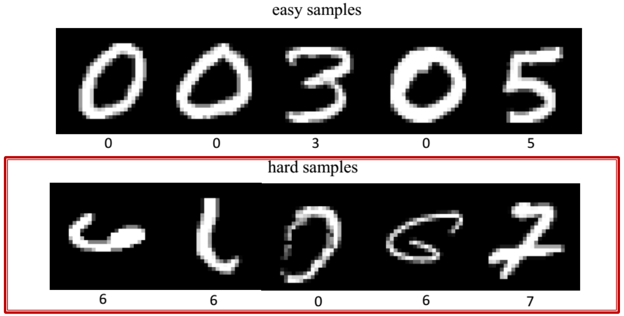

# Humans

- analyze the difficulty of each task

# Humans

- analyze the difficulty of each task
- accept tasks that are doable
  - such as classifying popular dog breeds

Dalmatian

Golden retriever

# Humans

- analyze the difficulty of each task
- accept tasks that are doable
  - such as classifying popular dog breeds
- but refuse tasks that are too hard
  - such as classifying exotic dog breeds
- they just say "sorry, can't do it"

Dalmatian

Golden retriever

Tibetan Mastiff

Peruvian Inca Orchid

# Motivation

- we refer to this as realism
- this trait is critical for many applications,
  - circumventing risk in autonomous driving

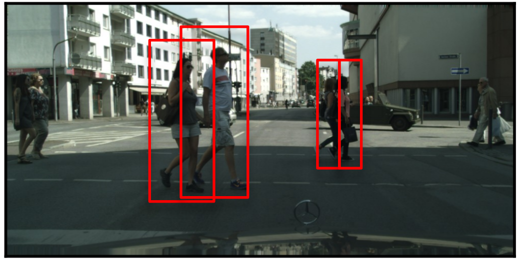

missed
detections

# Motivation

- we refer to this as realism
- this trait is critical for many applications,
  - circumventing risk in autonomous driving
  - tumor detection and classification

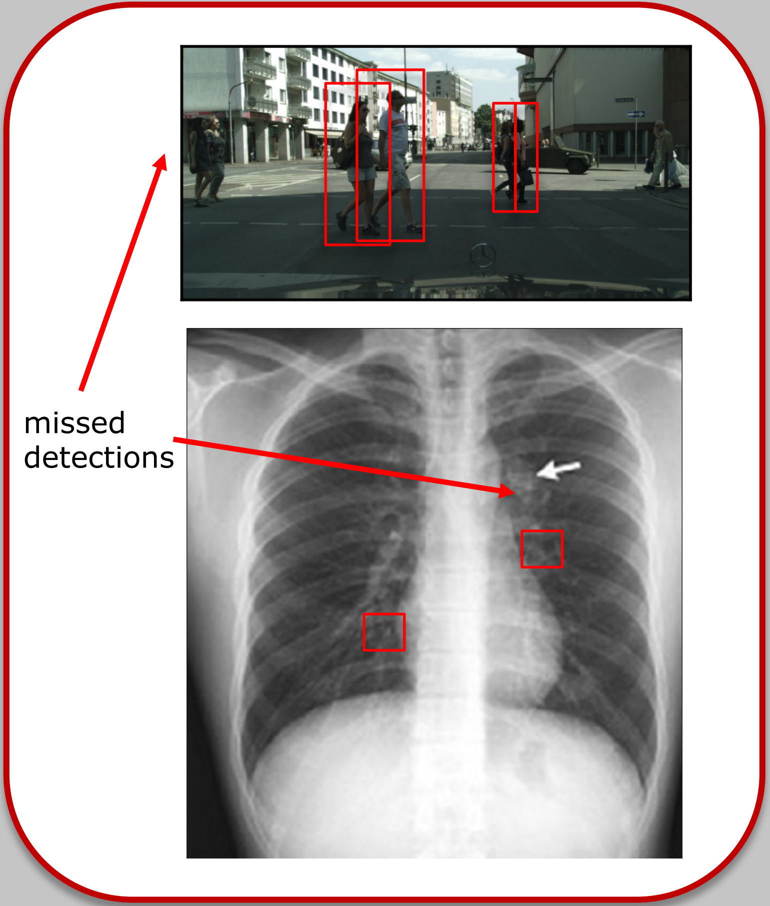

# Motivation

- we refer to this as realism
- this trait is critical for many applications,
  - circumventing risk in autonomous driving
  - tumor detection and classification
- vision system should either
  - reject to perform the hard task
  - request additional information from other sensors

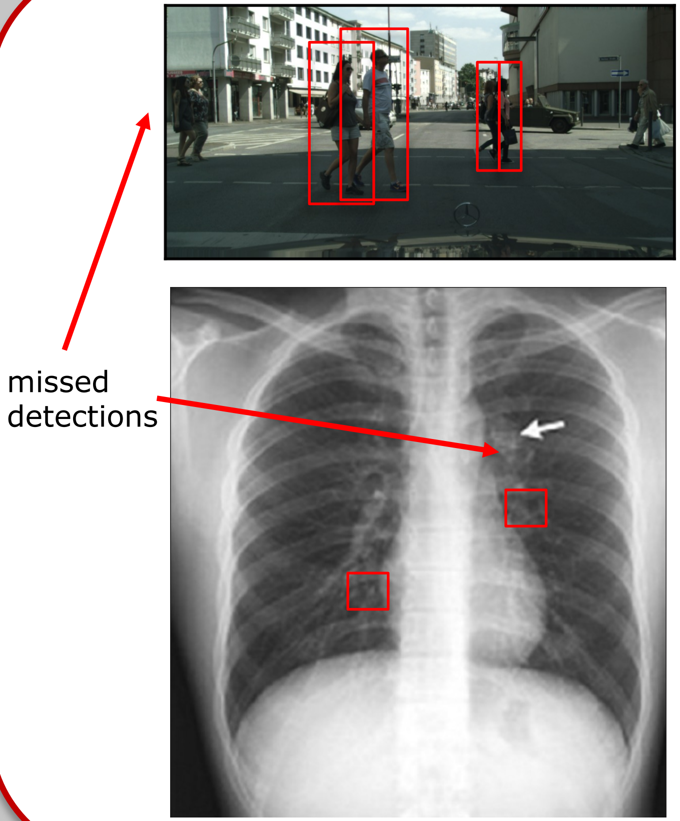

# Realistic classifier

- standard classifier
  - classifies all examples
- realistic classifier is defined as a classifier
  - that rejects examples deemed too hard
  - to guarantee a target performance on the accepted examples

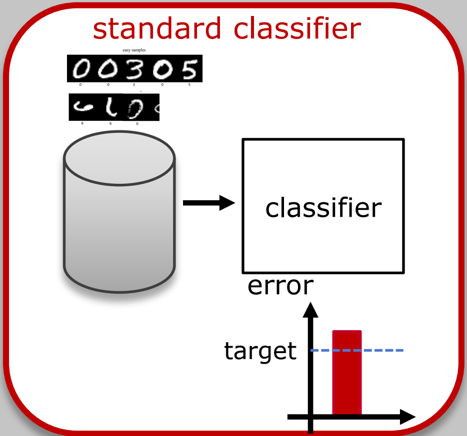


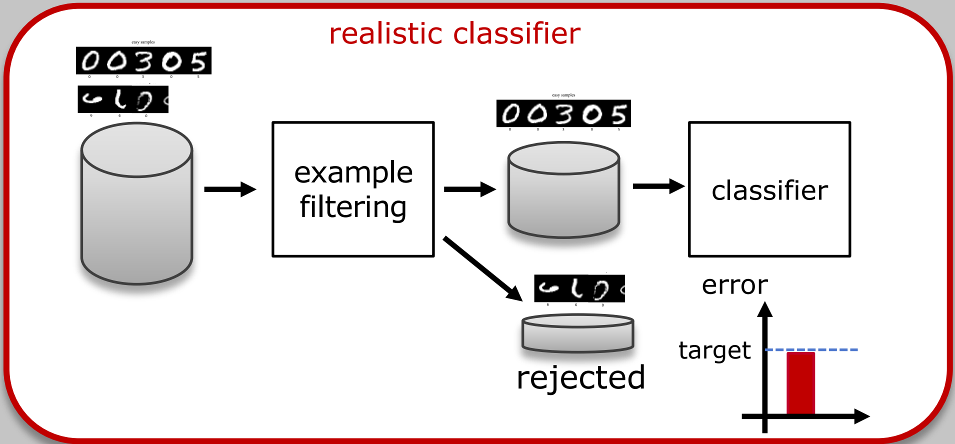

# Realistic classifier

- implemented as a sequence of
  - hardness predictor: assigns hardness score to each example
  - score thresholded to reject hard examples
  - classifier only applied to examples that can be "realistically" classified

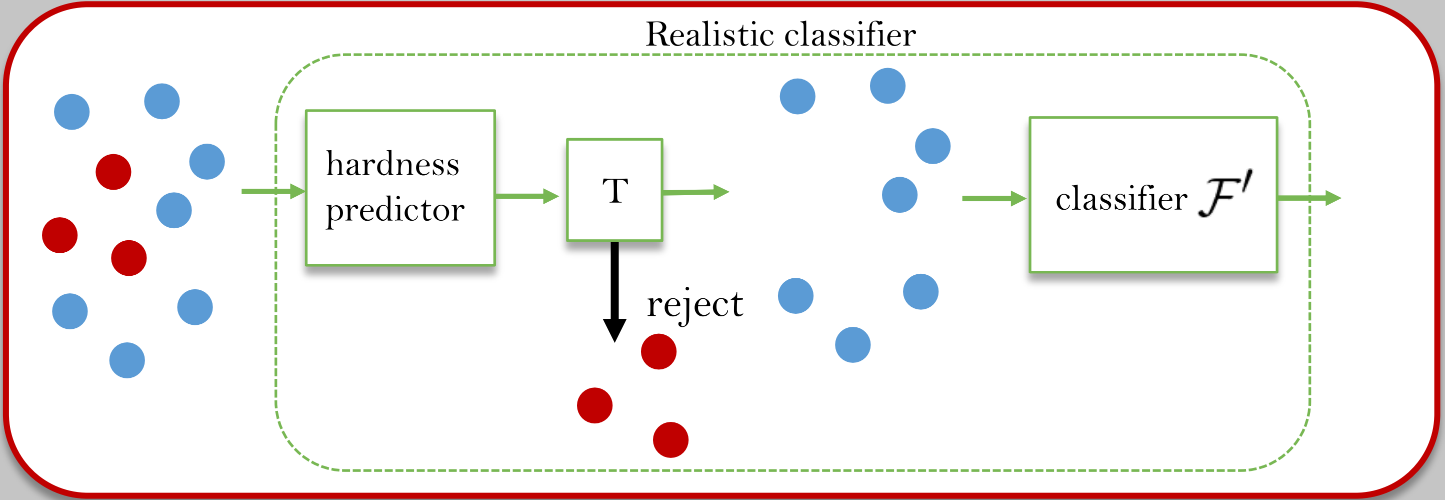

# Learning Hardness Predictor

- adversarial learning procedure
  - classifier and hardness predictor learned alternately
  - classifier learned with variant of cross entropy loss $L_m$
  - hardness predictor learned with a loss that encourages large hardness scores for misclassified examples

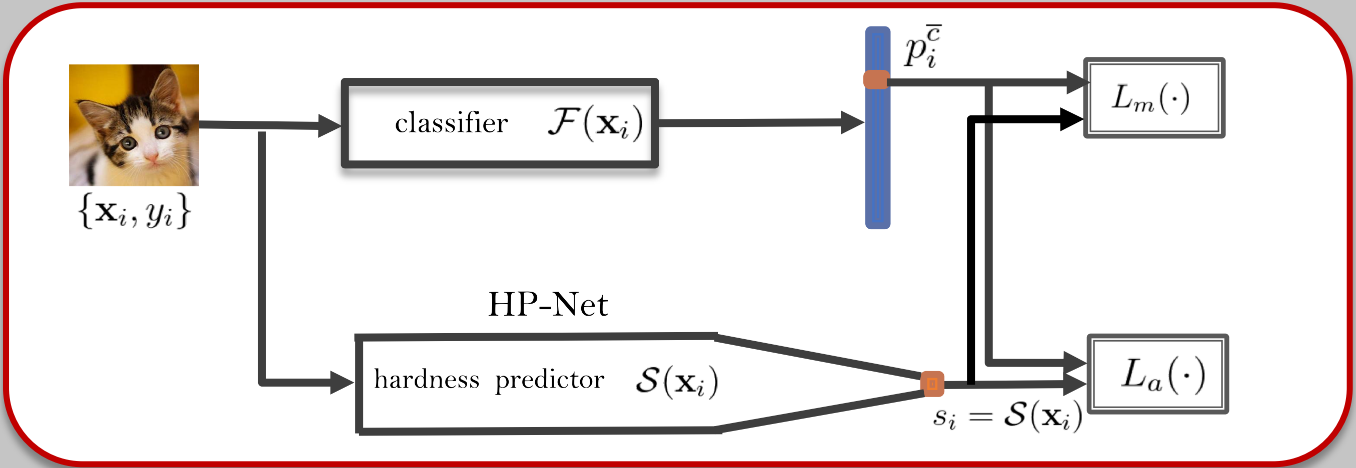

# Experiments

- Benefits of joint optimization:
  - accuracy and hardness score distribution evolution on MNIST

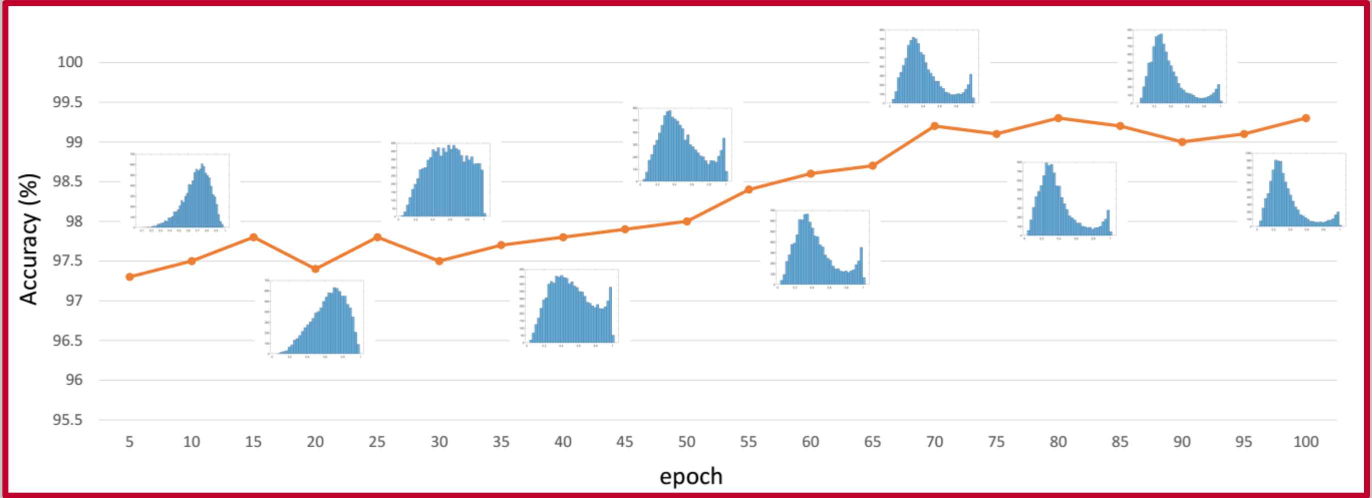


- as the classifier improves the distribution shifts to the left
- fewer hard examples, consistent with better classifier

# Experiments

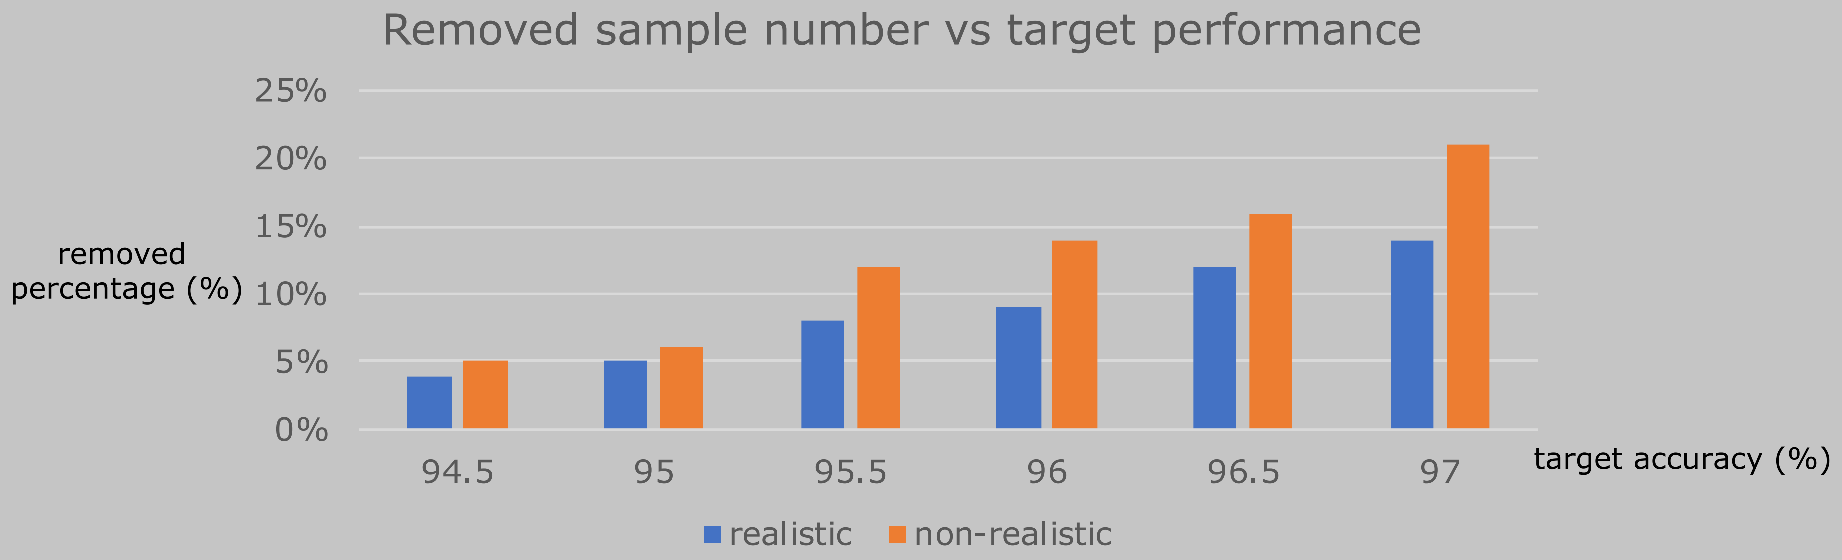


- In order to guarantee a target performance, realistic predictors can accept and classify more examples than non-realistic ones.

# Conclusions

- computer vision systems try to process all instances
- this optimistic attitude can lead to critical failures in some applications
- realistic classifiers reject some examples to guarantee a target performance on the ones they process
- proposed an adversarial architecture for realistic prediction, based on joint learning of hardness predictor and classifier
- this was shown to
  - improve classification when hard examples are rejected
  - superior to thresholding of confidence scores for example rejection

# Thank You!

Welcome to our poster for more details