Learning with Limited Supervision

Alessandro Giusti

Dalle Molle Institute for Artificial Intelligence

USI and SUPSI
Lugano, Switzerland

Scuola universitaria professionale
della Svizzera italiana
SUPSI

IDSIA

Contact: alessandrog@idsia.ch

https://idsia-robotics.github.io/

Dalle Molle Institute for AI (IDSIA)

Lugano, Switzerland

USI/SUPSI
IDSIA

Dalle Molle Institute for AI (IDSIA)

Lugano, Switzerland

Dalle Molle Institute for AI (IDSIA)

A world-class research institute on AI founded in 1988 in Lugano

120+ staff working on:

- Deep learning
- Statistics and data mining
- Visual Computing
- Autonomous robotics
- Natural language processing
- Operations research

Dalle Molle Institute for AI (IDSIA)

A world-class research institute on AI founded in 1988 in Lugano

120+ staff working on:

- Deep learning
- Statistics and data mining
- Visual Computing
- **Autonomous robotics**
- Natural language processing
- Operations research

GTC keynote on Tuesday April 5, 2016
Jensen H Huang, CEO of NVIDIA

Robotics lab: People & Robots

People

- 5 Postdoc researchers
- 2 Researchers without PhD
- 3 PhD Students

Robots

- 6 ground (1-10 kg)
- 10+ flying (20-500 g)
- 30+ swarm (50-200 g)

90 sqm **robot lab** equipped with 18-camera mocap system

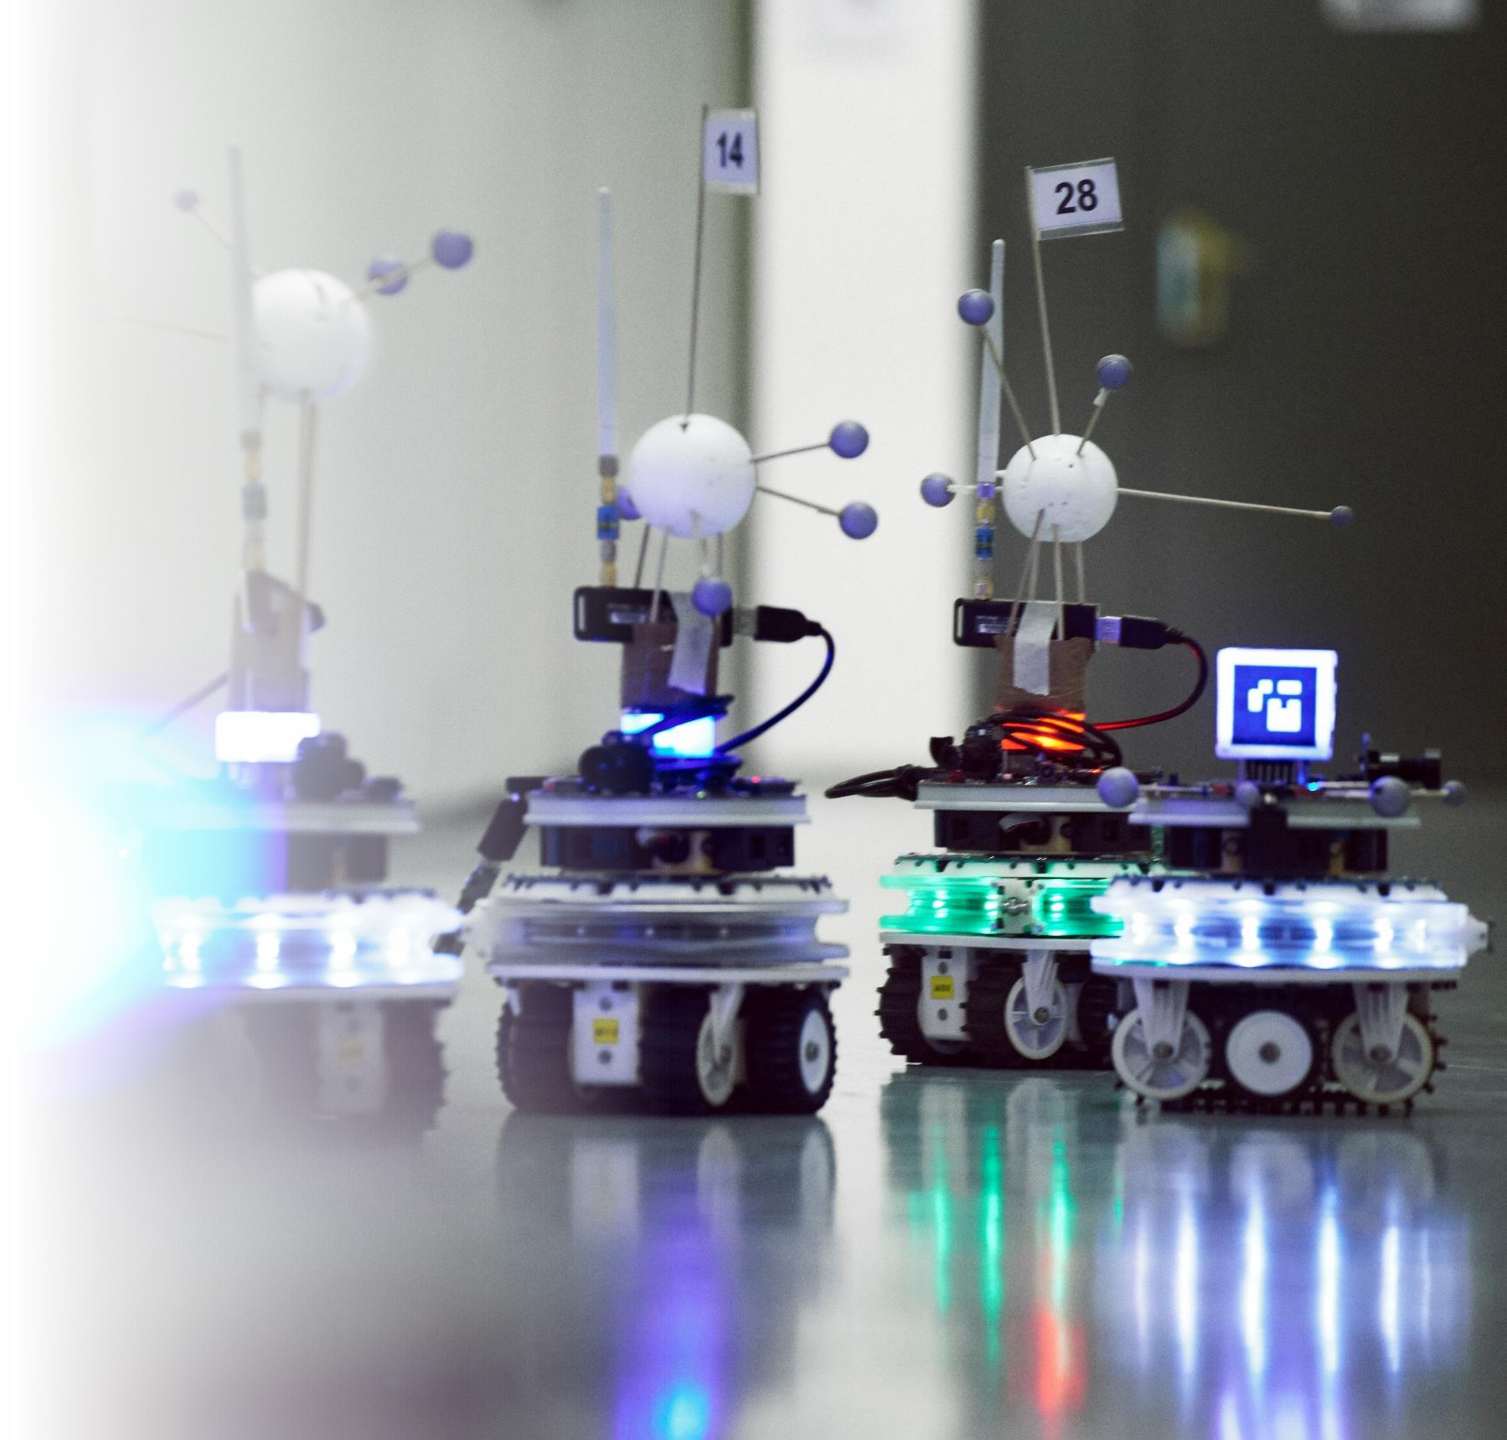

Research Themes

Perception for Autonomous Mobile Robots

- Self-supervised Deep Learning
- Nano-quadrotors

Human-Robot Interaction

Robot Swarms

Industrial applications

- Manufacturing / Machining
- Optical Inspection
- Space

Visual control of drones in forests

Finalist,
AAAI video awards,
2016, Phoenix, USA

Giusti et al., Robotics and Automation Letters, 2016 -- https://youtu.be/umRdt3zGgpU

Visual control of nano drones in human proximity

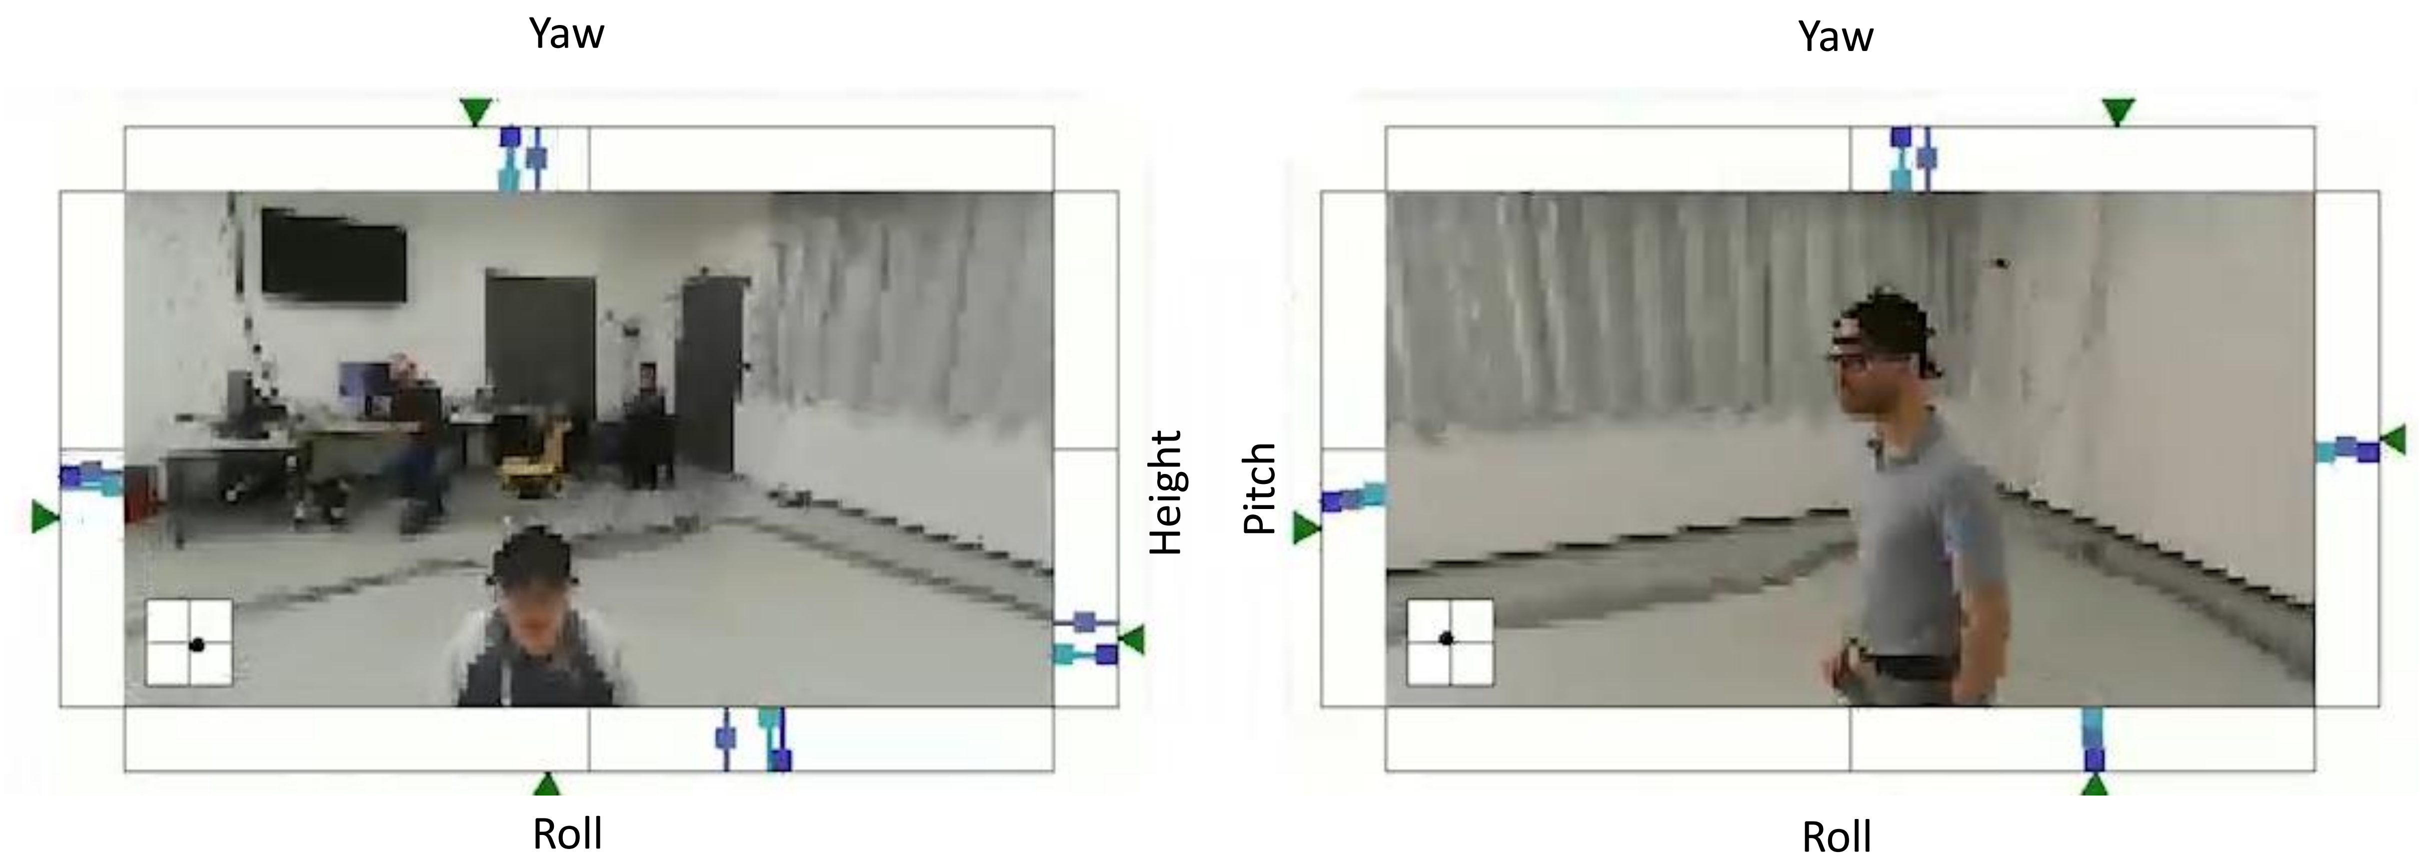

▲ ground truth control

■ predicted values

Mategazza et al., ICRA 2019 -- https://github.com/idsia-robotics/proximity-quadrotor-learning

Visual control of nano drones in human proximity

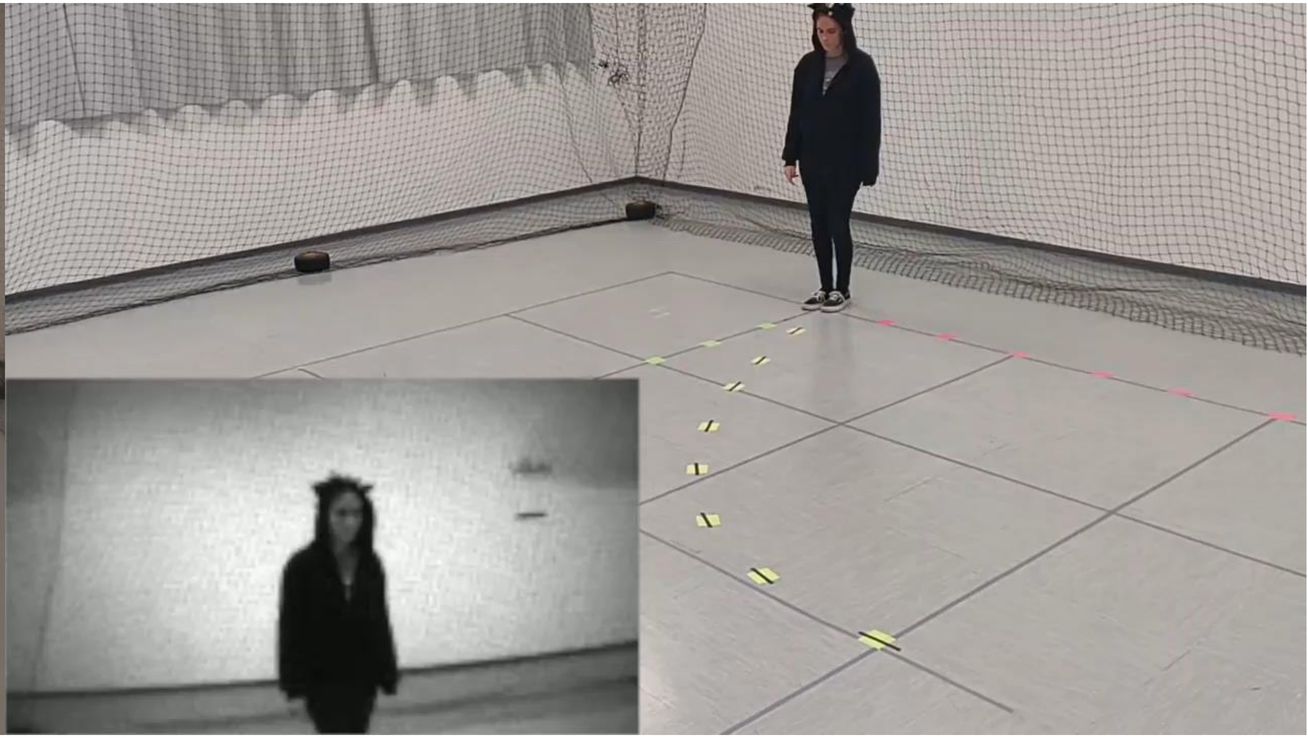

- Full onboard CNN inference and control at **130 frames per second on a 30-gram drone**
- **< 1 mJ per frame** on Parallel Ultra Low Power platform
- **Self-supervised learning** approach improves performance without explicit ground truth

Palossi et al., Internet of Things Journal, 2021 -- https://github.com/idsia-robotics/pulp-frontnet

Optical analysis of circuit boards

DELVITECH® Artificial Intelligence Applied to AOI

Deep Neural Networks **predict multiple component characteristics** (type, number of pins, material, dimensions, …) for self-programming of Automated Optical Inspection machines

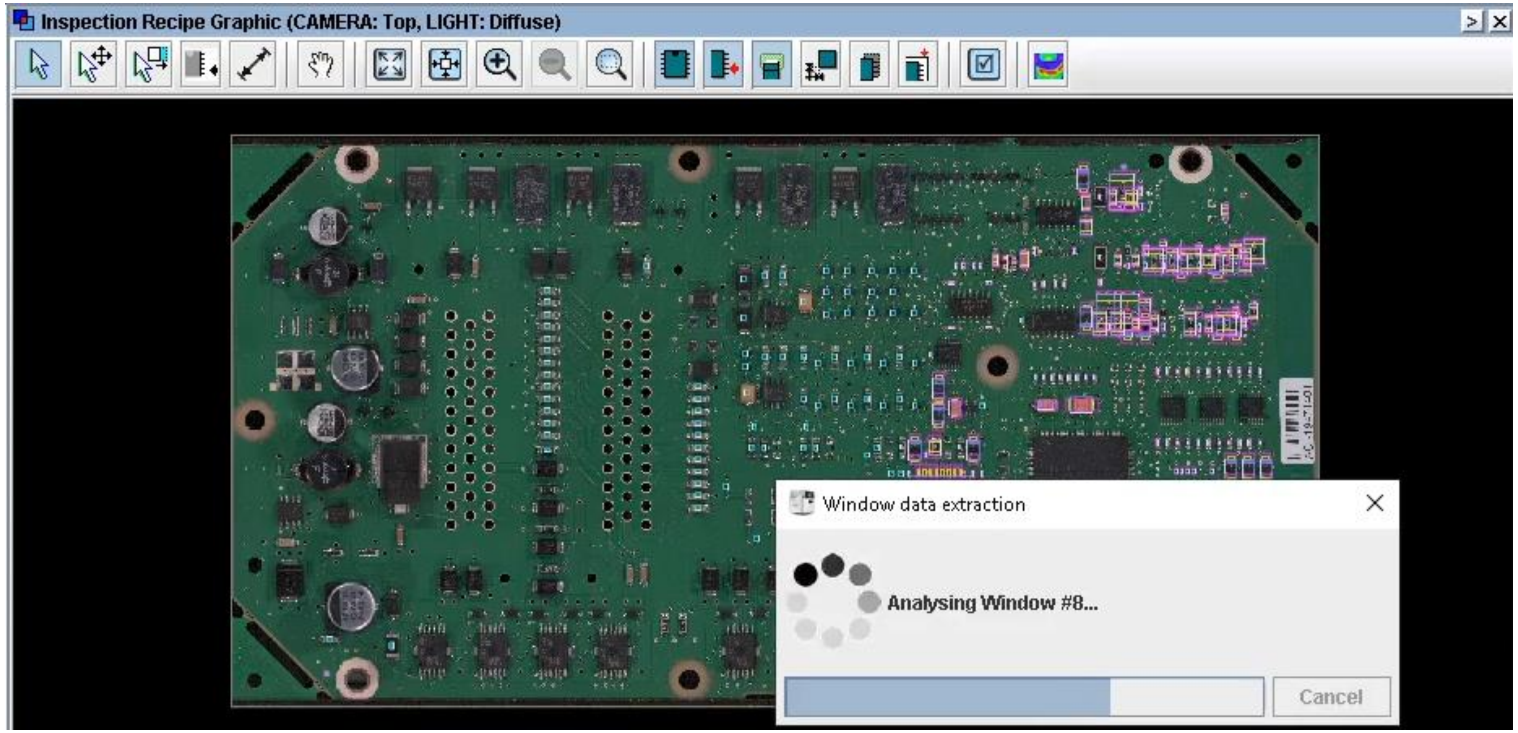

Deep Neural Networks **trained without ground truth** estimate component precise localization and rotation

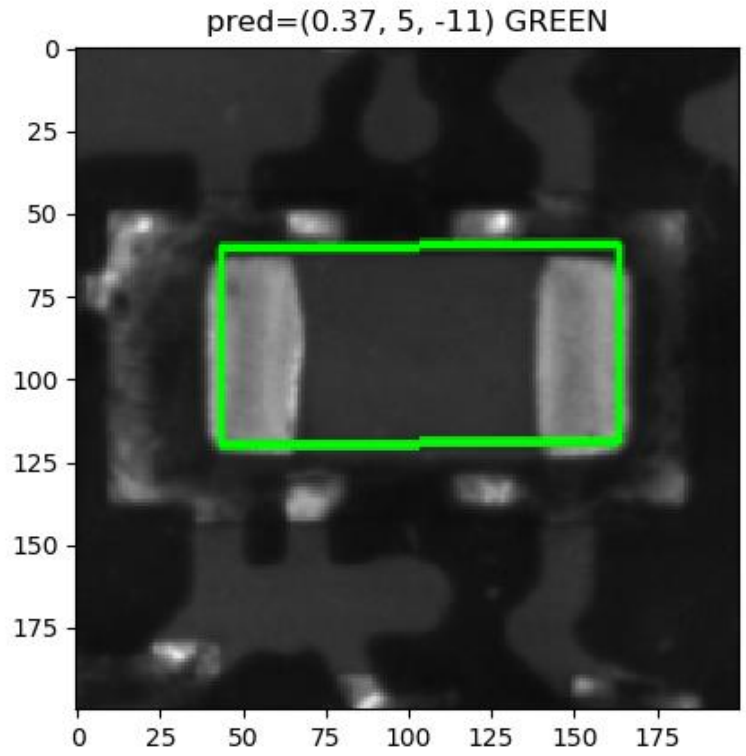

with Delvitech SA

Innosuisse

Visual detection of counterfeit brand labels

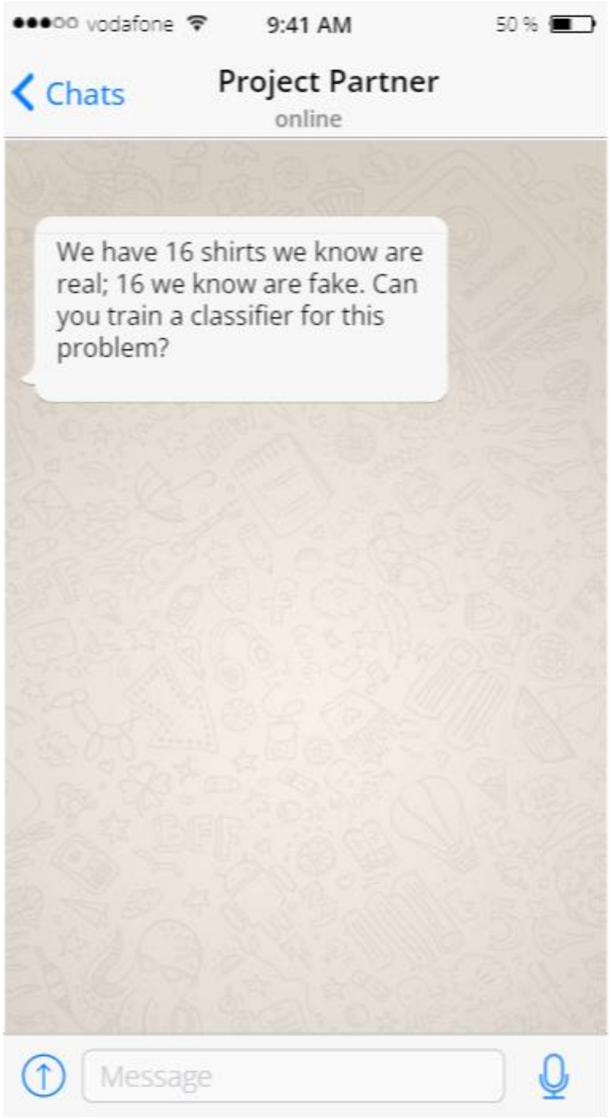

Training data augmentation to the rescue

Original

6 out of 500 random variations

Training data augmentation

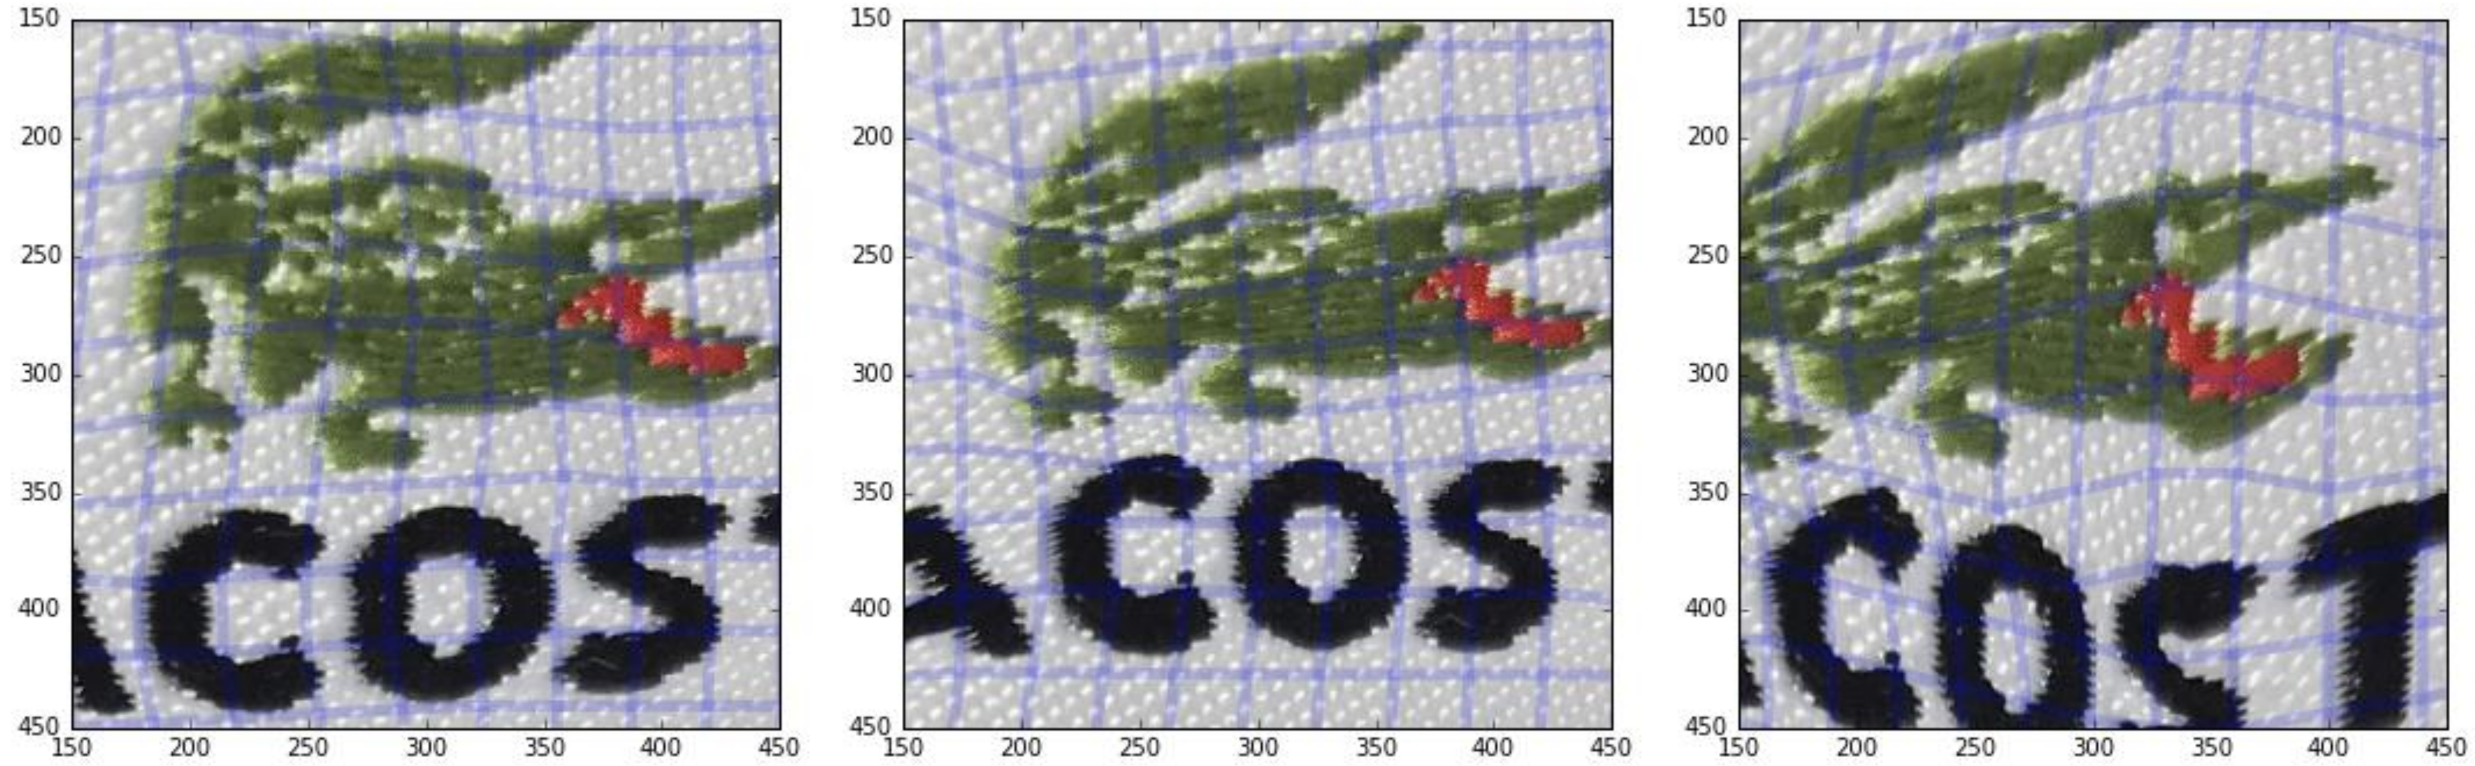

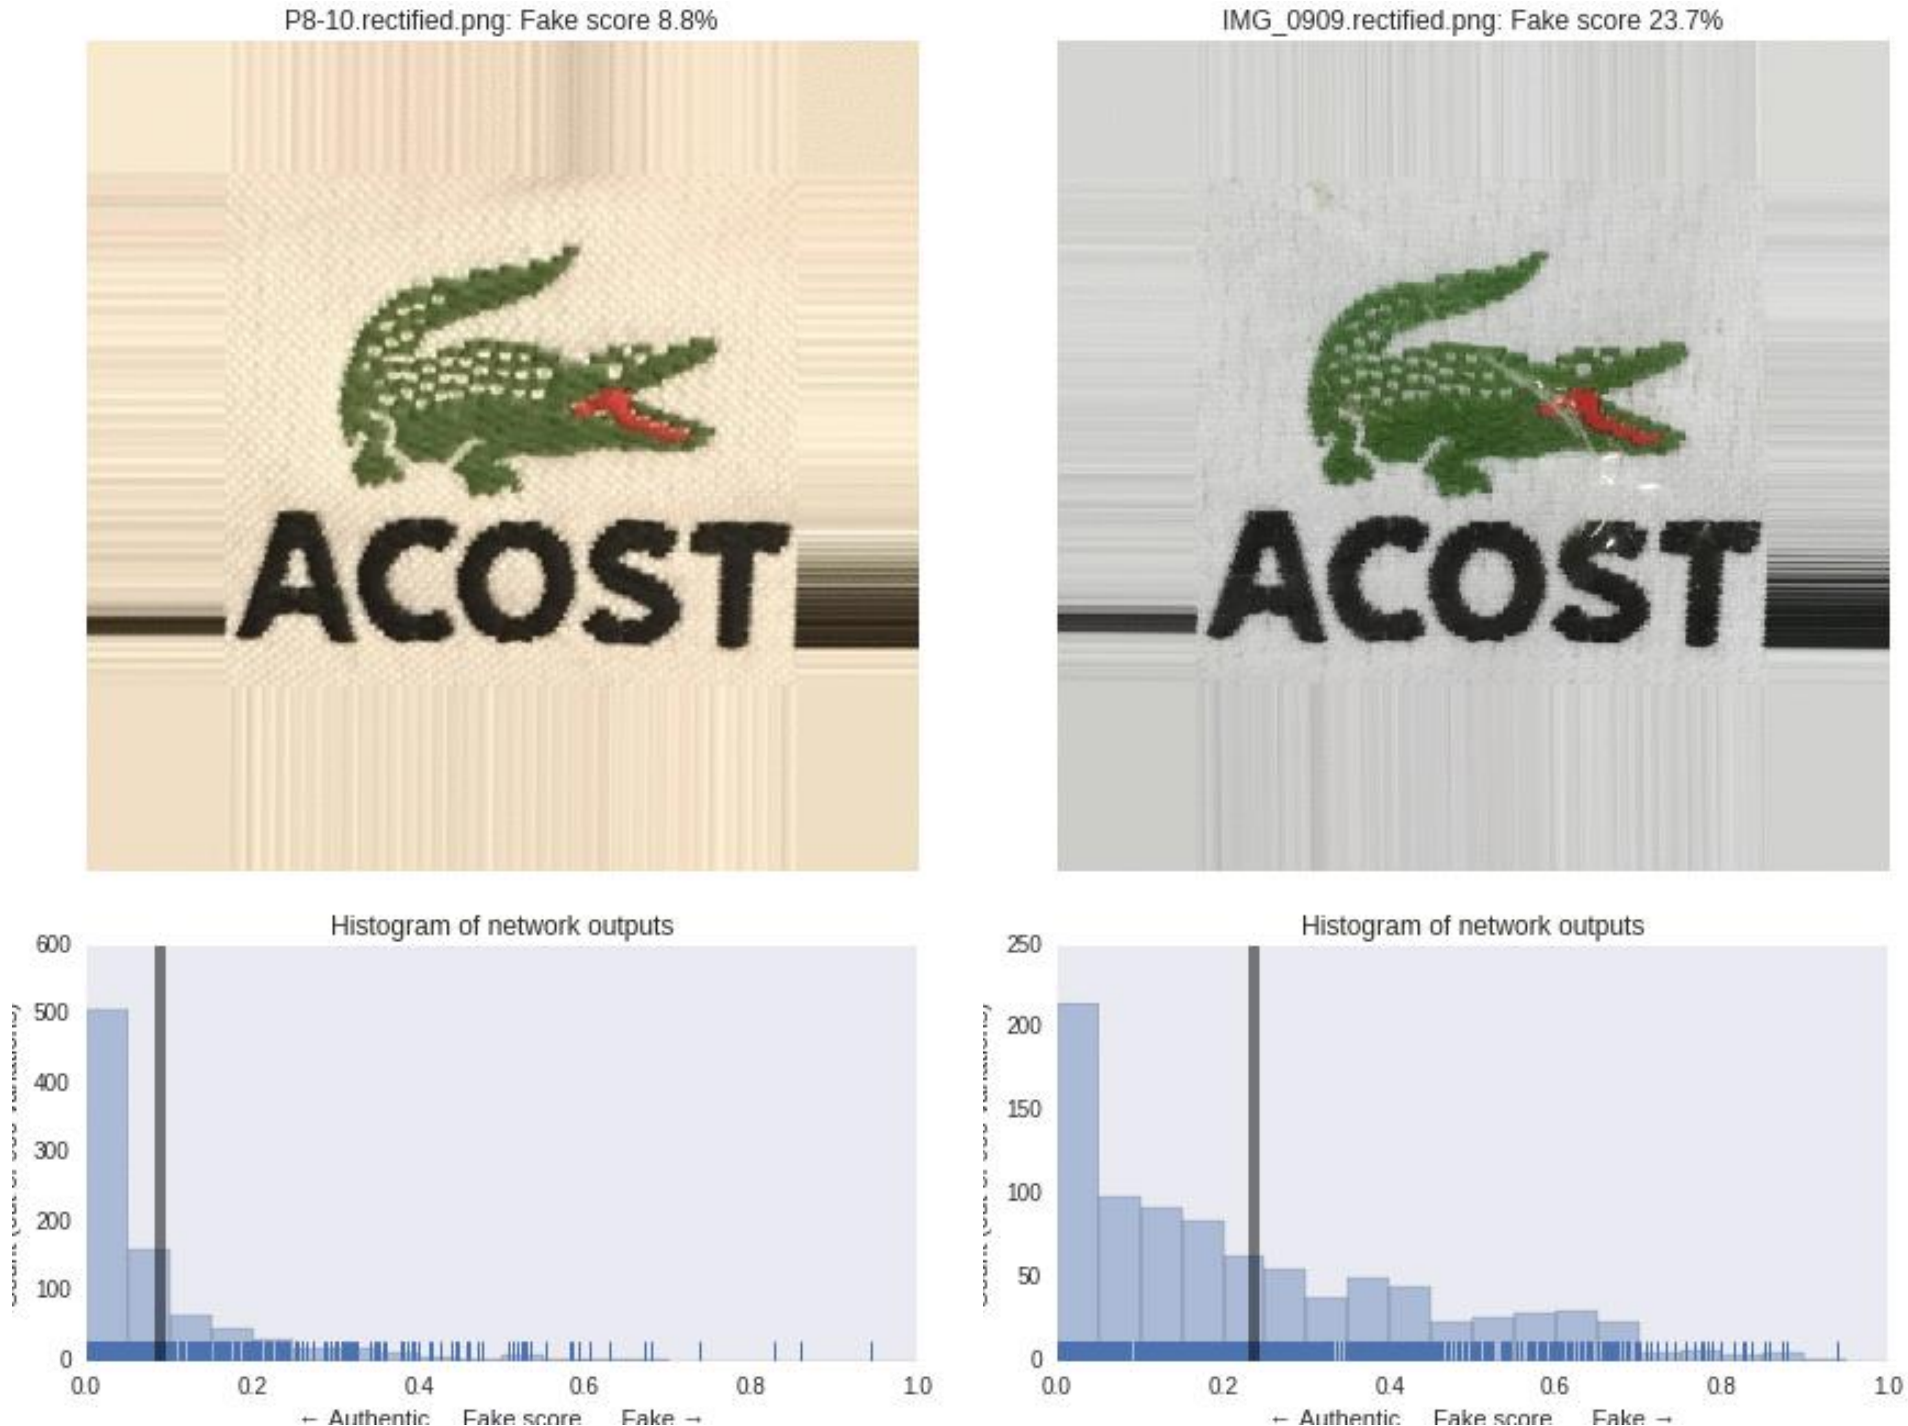
P8-10.rectified.png: Fake score 8.8%
IMG_0909.rectified.png: Fake score 23.7%
ACOST
ACOST
Histogram of network outputs
Histogram of network outputs
0
100
200
300
400
500
600
0
50
100
150
200
250
0.0
0.2
0.4
0.6
0.8
1.0
← Authentic Fake score Fake →
← Authentic Fake score Fake →

Training data is RARE!

For most real-world ML problems

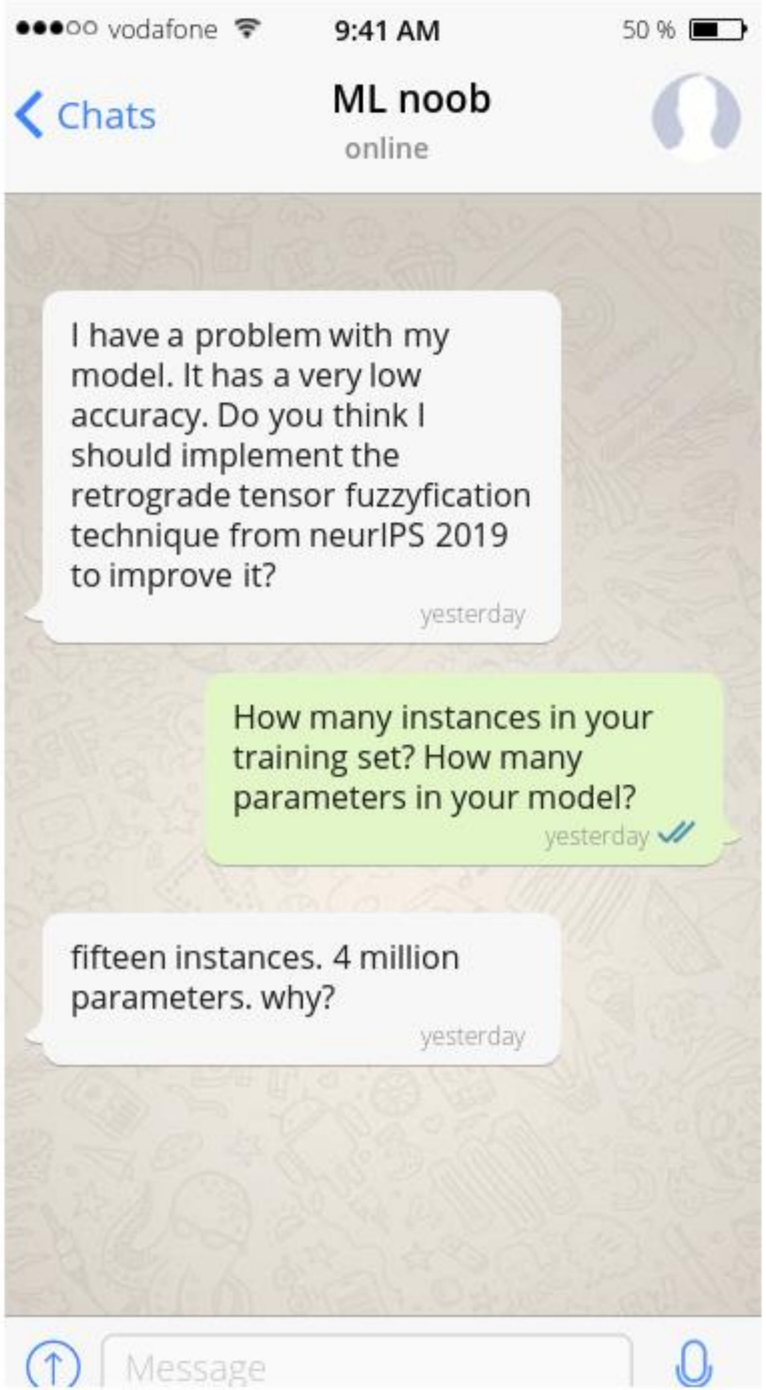

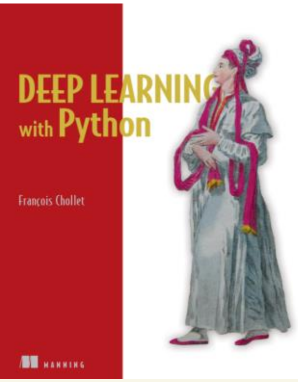

The ML Pipeline (F. Chollet)

- Define the problem at hand and the data on which you'll train. Collect this data, or annotate it with labels if need be.
- Choose how you'll measure success on your problem. Which metrics will you monitor on your validation data?
- Determine your evaluation protocol: hold-out validation? K-fold validation? Which portion of the data should you use for validation?
- Develop a first model that does better than a basic baseline: a model with statistical power.
- Develop a model that overfits.
- Regularize your model and tune its hyperparameters, based on performance on the validation data. A lot of machine-learning research tends to focus only on this step—but keep the big picture in mind.

Plan of the lecture

- Part **1**: introduction
- Part **2**: warm-up on the CIFAR-10 dataset
- Part **3**: what is self-supervised learning?
- Part **4**: implement&test a simple self-supervised learning method
- Part **5**: some examples of self-supervised learning in robotics

Plan of the lecture

- Part **1**: introduction
- **Part 2: warm-up on the CIFAR-10 dataset**
- Part **3**: what is self-supervised learning?
- Part **4**: implement&test a simple self-supervised learning method
- Part **5**: some examples of self-supervised learning in robotics