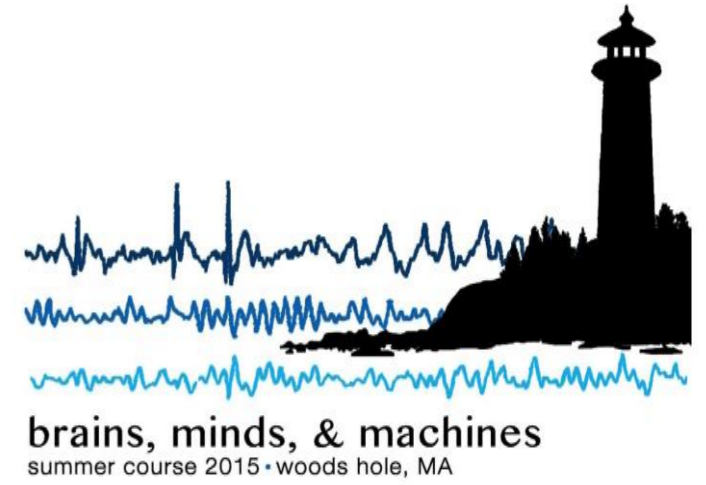


© RobotCub Consortium. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Giulia Pasquale
PhD student



IIT, iCub Facility –
University of Genoa, DIBRIS –
Laboratory for Computational and Statistical Learning –



Picture of Lorenzo Natale removed due to copyright restrictions. Please see the video.

Picture of Lorenzo Rosasco removed due to copyright restrictions. Please see the video.

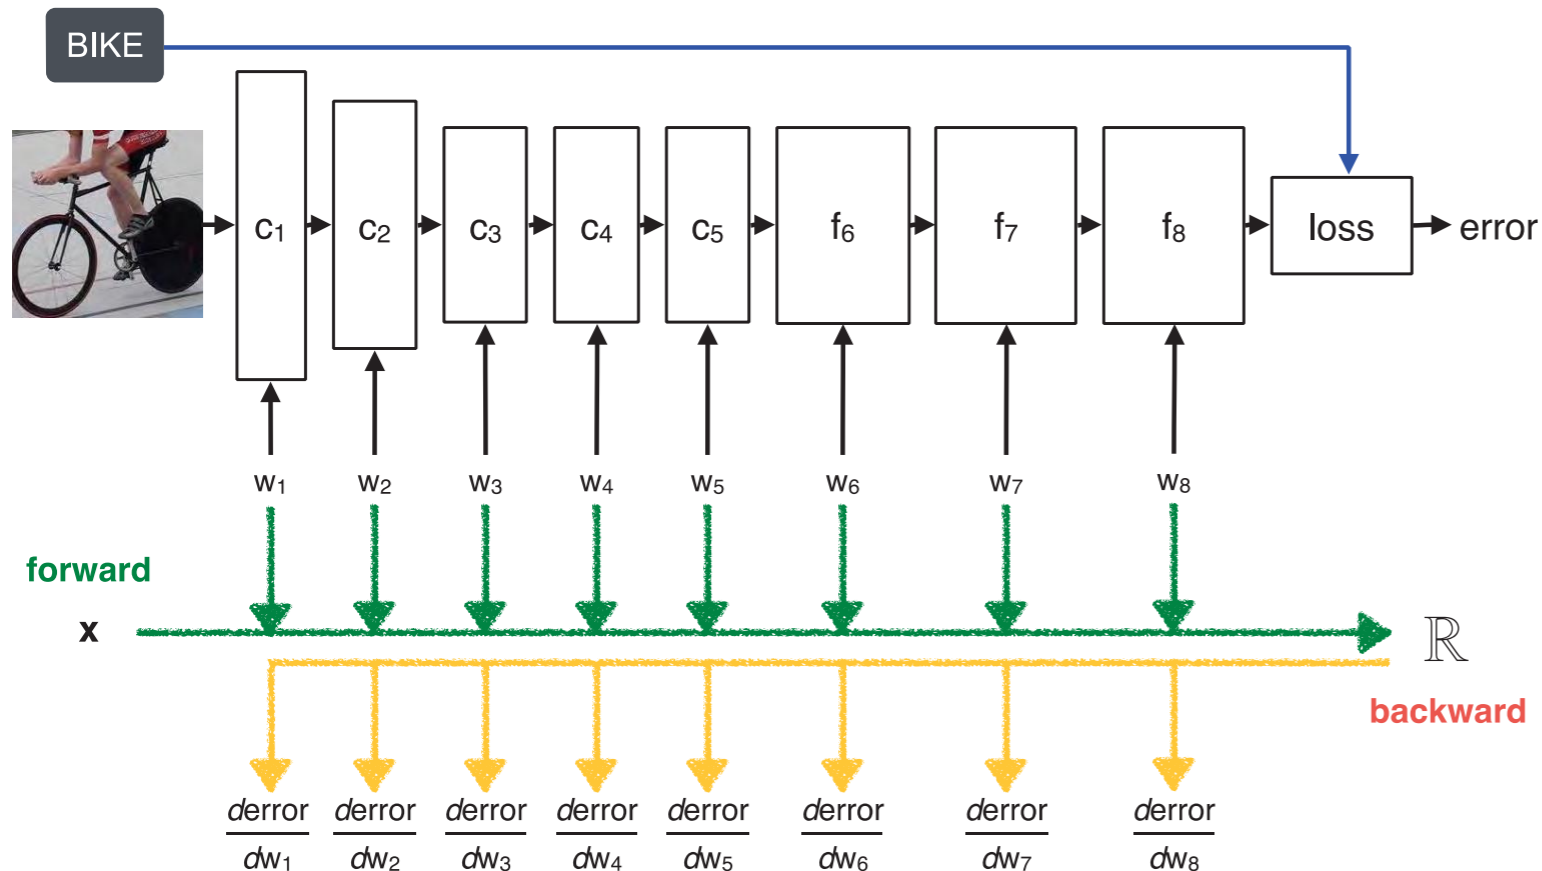
Supervisors and collaborators

Picture of Carlo Ciliberto removed due to copyright restrictions. Please see the video.

Picture of Francesca Odone removed due to copyright restrictions. Please see the video.

Deep Learning Breakthrough in Computer Vision

DEEP NETWORKS



BIG DATASETS

IMAGENET15M

[Deng et al. '09]

Number of Labeled Images

SUN, **131K**
[Xiao et al. '10]

LabelMe, **37K**
[Russell et al. '07]

PASCAL VOC, **30K**
[Everingham et al. '06-'12]

Caltech101, **9K**
[Fei-Fei, Fergus, Perona, '03]

Credits: A. Vedaldi

© Oxford Visual Geometry Group. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Credits: Fei-Fei Li

© Source Unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.


Deep Learning Breakthrough in Computer Vision

Figure removed due to copyright restrictions. Please see the video.


Source: Figures 9, 11 & 12 from Russakovsky, Olga, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang et al. "Imagenet large scale visual recognition challenge." *International Journal of Computer Vision* 115, no. 3 (2015): 211-252.

Deep Learning Breakthrough in Computer Vision

IMAGENET
PRE-TRAINING




Krizhevsky et al (2012)



Courtesy of Neural Information Processing Systems. Used with permission.
Source: Krizhevsky, Alex, IlyaSutskever, and Geoffrey E. Hinton. "Imagene classification with deep convolutional neural networks." In Advances in neural information processing systems, pp. 1097-1105. 2012.

PASCAL VOC 2007




© Andrea Vedaldi. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Credits: A. Vedaldi

Meanwhile, in Robotics...

Image of a baby removed due to copyright restrictions. Please see the video.

Meanwhile, in Robotics...



© RobotCub Consortium. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.



© DARPA. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.



© AUVSI. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.



© DARPA. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

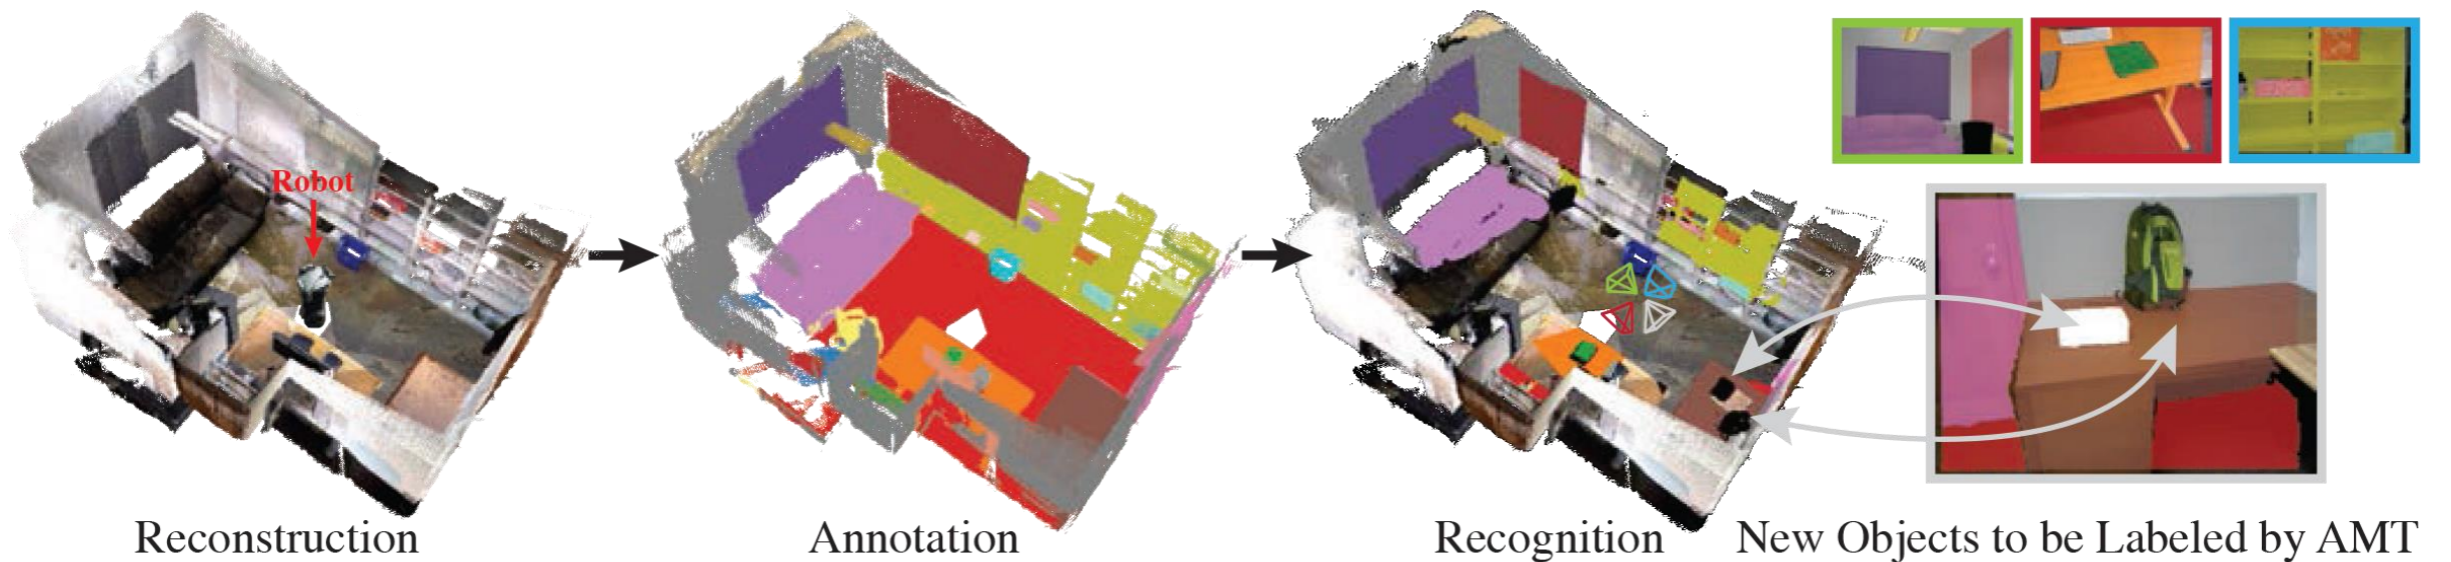
Meanwhile, in Robotics...

TELE-OPERATION

Image removed due to copyright restrictions.
Please see the video.

Image removed due to copyright restrictions.
Please see the video.

3D MAPPING & STRONG SUPERVISION




Courtesy of Shuran Song, Linguang Zhang and Jianxiong Xiao. License CC BY.

Song et al (2015), arXiv: 1507.02703

Setting: Interactive Object Learning

Verbal
instructions
of a "teacher"

Robot's attention
(motion, color-based
segmentation)

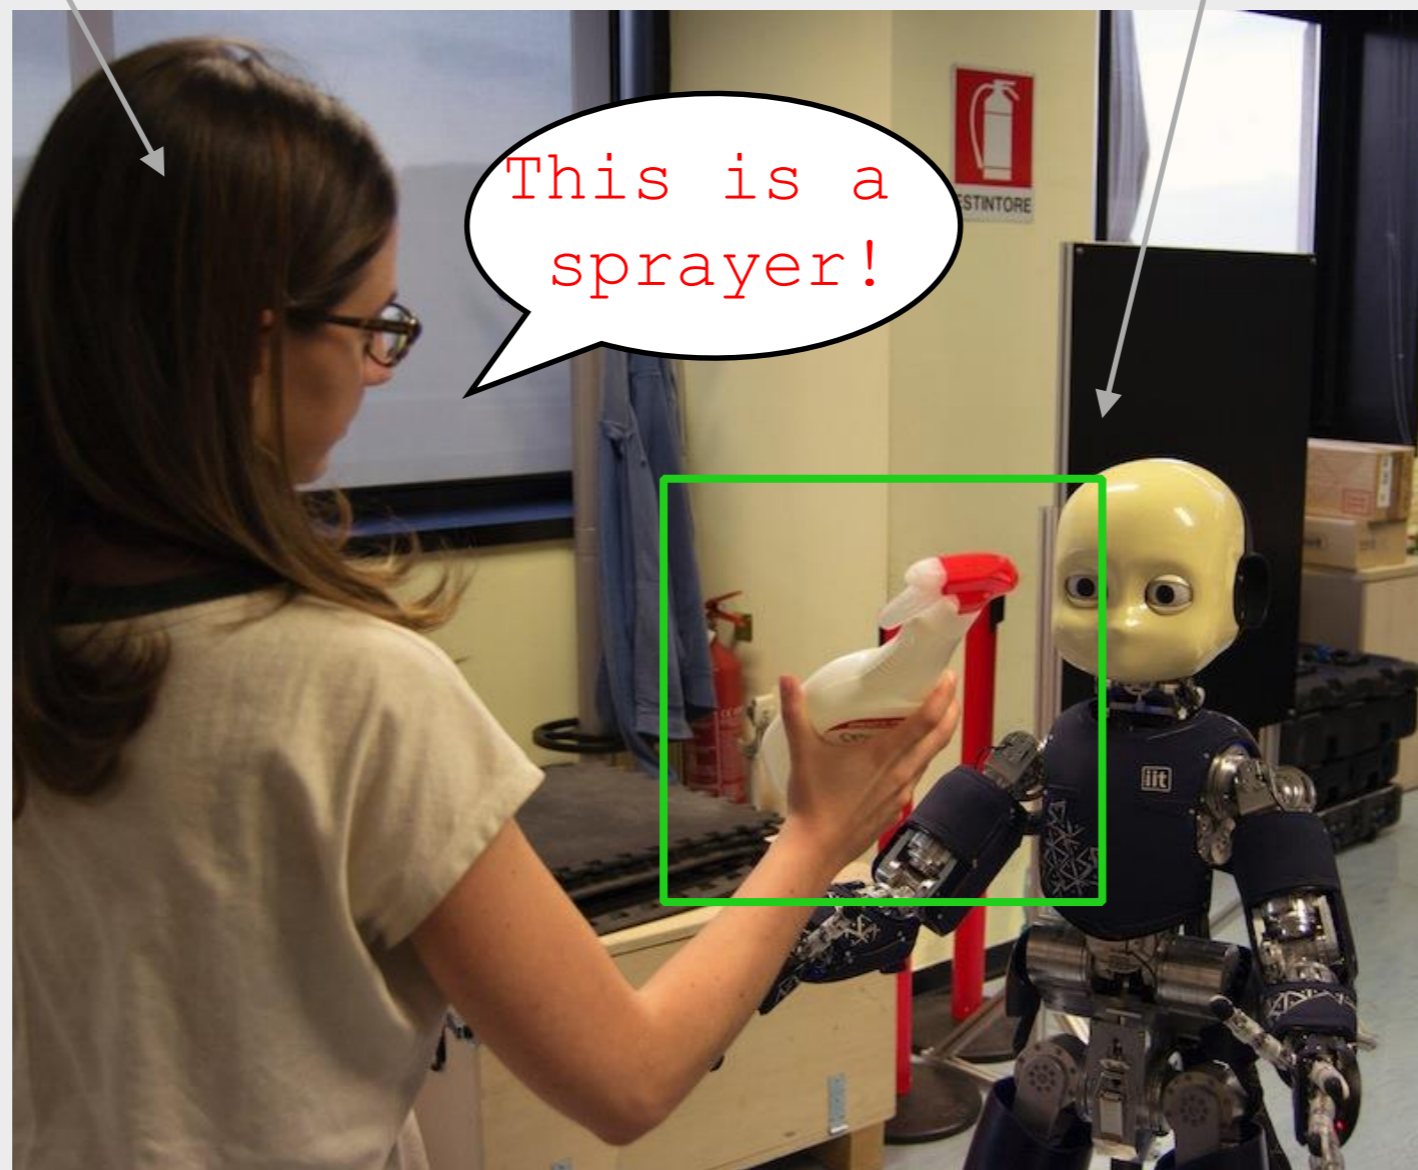


© RobotCub Consortium. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Setting: On the fly Recognition

Verbal
instructions
of a "teacher"

Robot's attention
(motion)



Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale. Used with permission.


Applications: Interactive Object Learning & On the fly Recognition

Verbal
Supervision

Segmentation

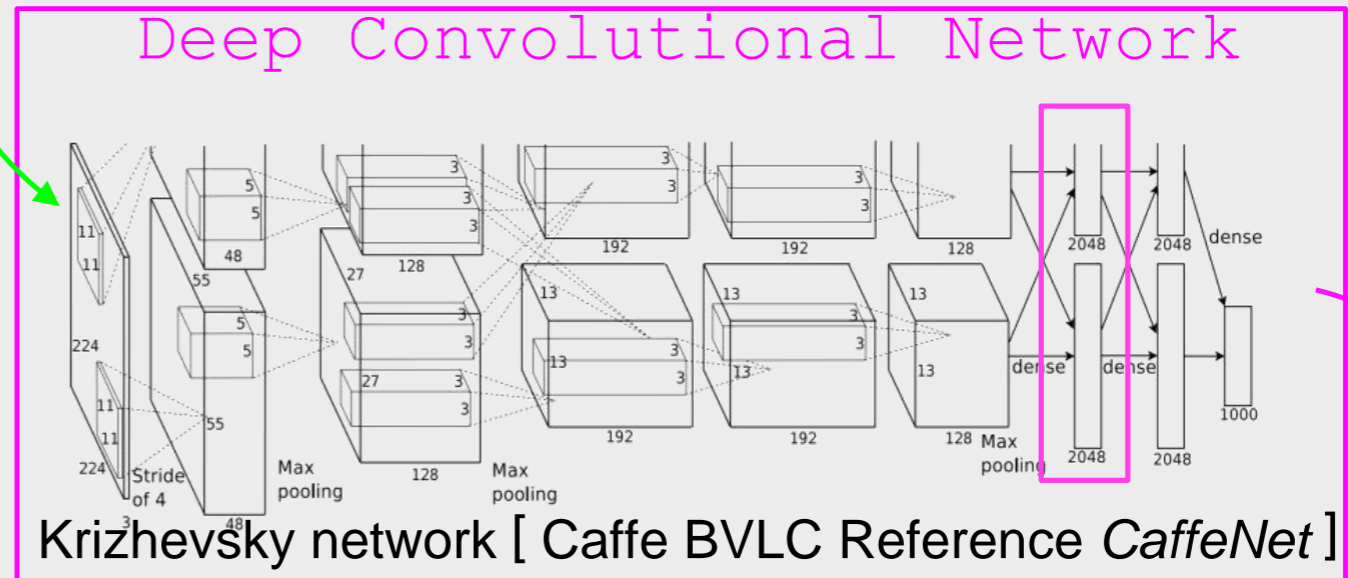
Representation
Extraction

Linear Classifier

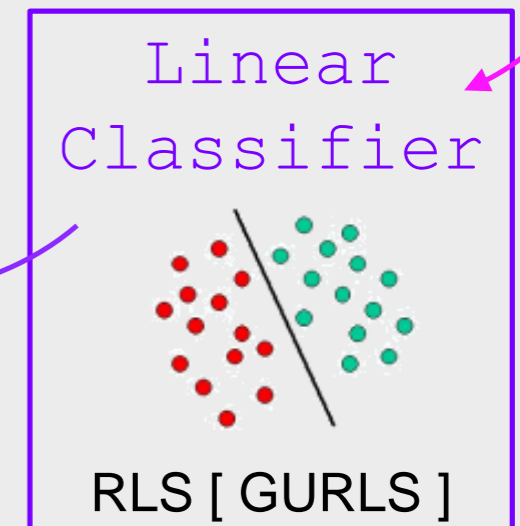



Motion,
Color & Luminance

© RobotCub Consortium. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.




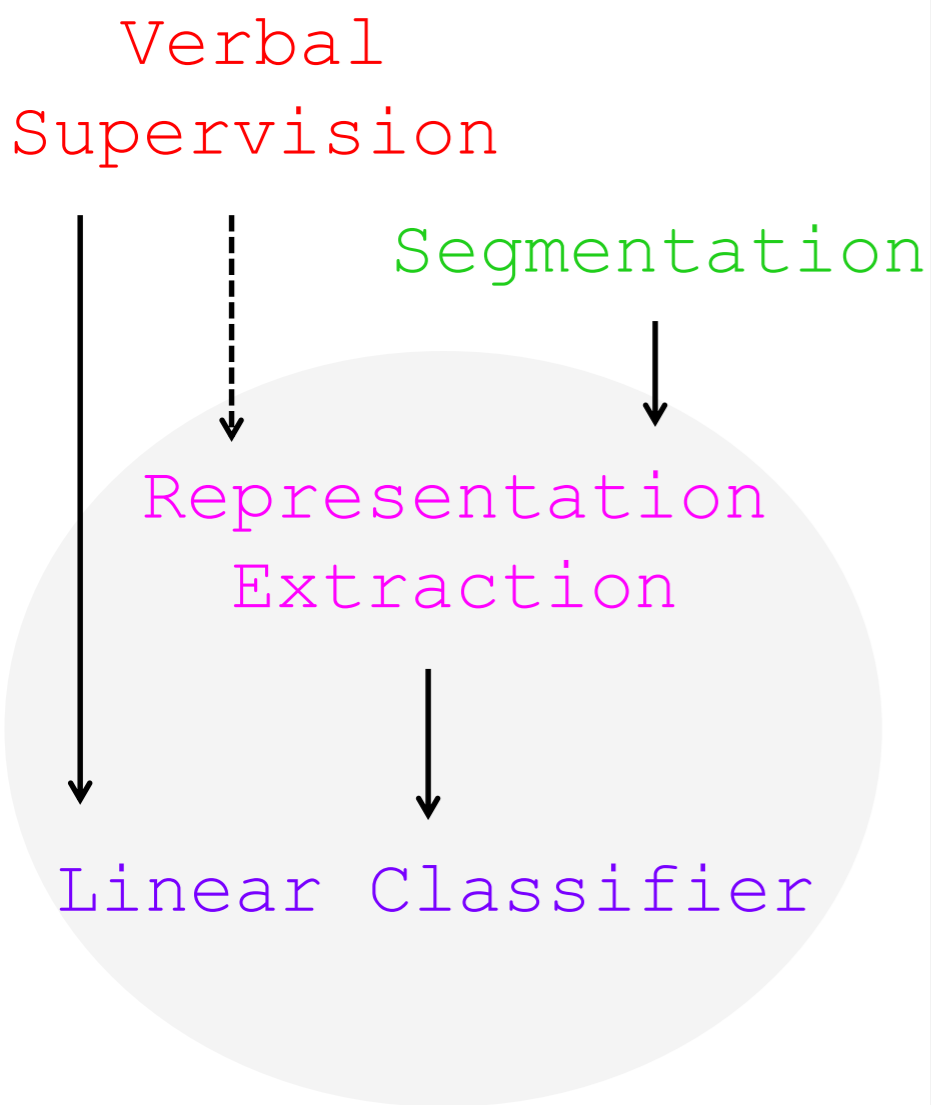
Courtesy of Neural Information Processing Systems. Used with permission.
Source: Krizhevsky, Alex, IlyaSutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." In Advances in neural information processing systems, pp. 1097-1105. 2012.



scores

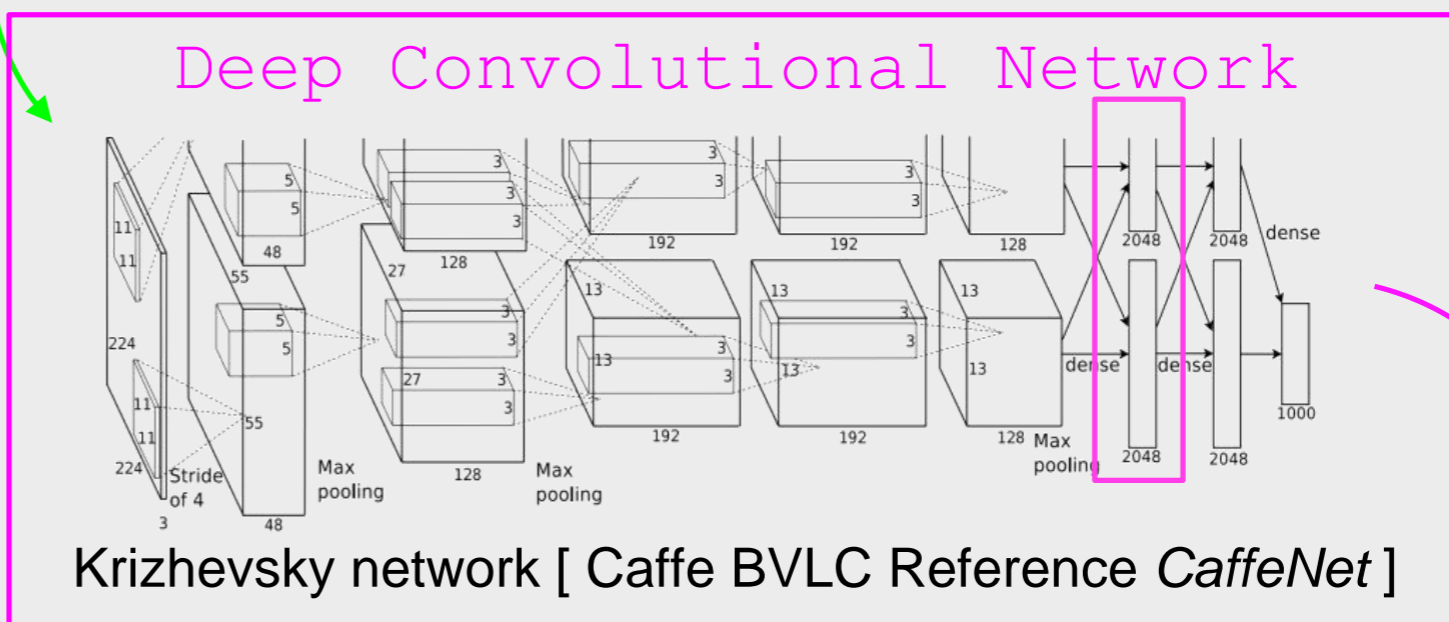
code

Applications: Interactive Object Learning & On the fly Recognition

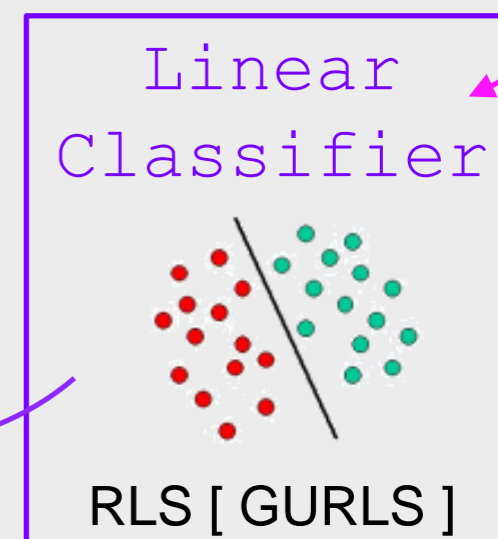



Motion,
Color & Luminance

Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco, Lorenzo Natale. Used with permission.
Source: Pasquale, Giulia, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco, Lorenzo Natale. "Teaching iCub to recognize objects using deep Convolutional Neural Networks." In MLIS@ICML, pp. 21-25. 2015.



Courtesy of Neural Information Processing Systems. Used with permission.
Source: Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." In Advances in neural information processing systems, pp. 1097-1105. 2012.



scores

An ideal robotic visual recognition system

- ✓ Self-supervised
- ✓ Reliable
- ✓ Exploits contextual information
- ✓ Learns incrementally

Teaching through time...

day 1 day 2 day 3 ... day X

What is this?


mug
mug
mug
mug
mug

detergent
detergent
detergent
detergent
detergent

sponge
sponge
sponge
sponge
sponge

© Source Unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Application: On the fly Recognition



Motion,
Color & Luminance

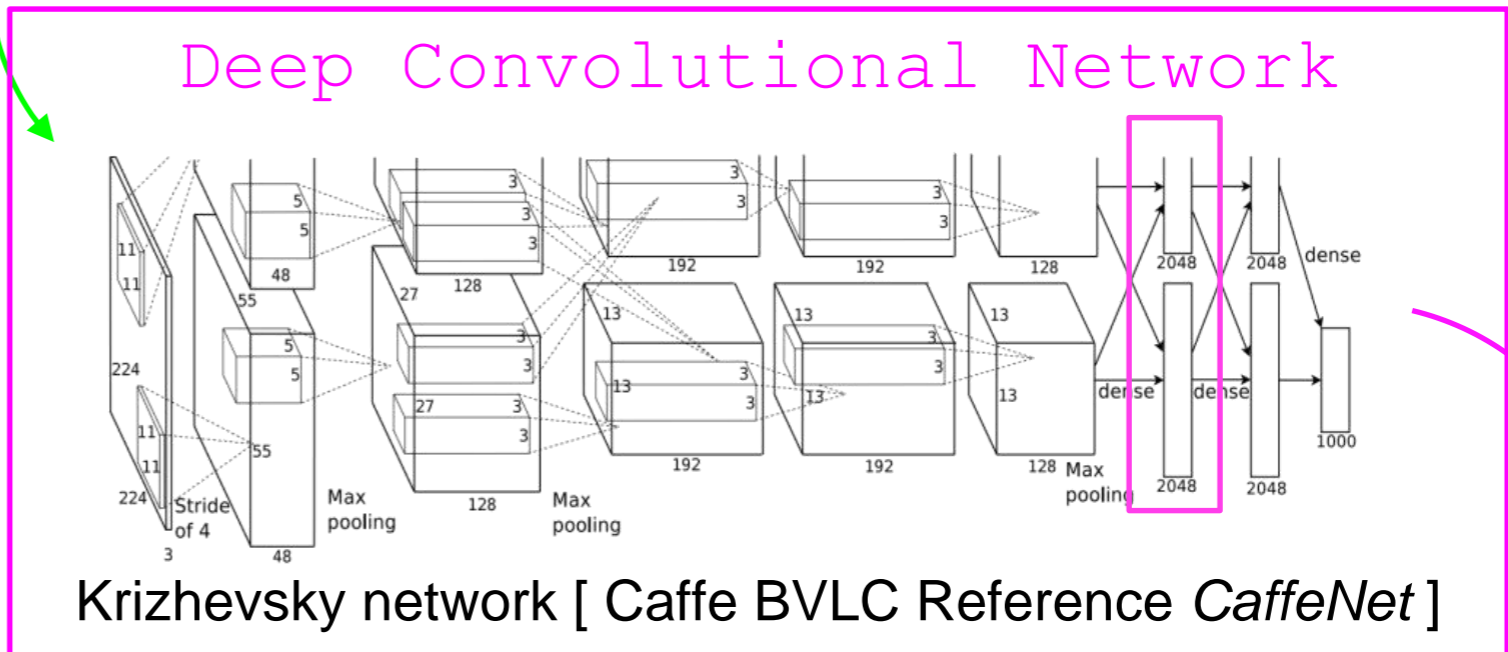
Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco, Lorenzo Natale. Used with permission.
Source: Pasquale, Giulia, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco, Lorenzo Natale. "Teaching iCubto recognize objects using deep Convolutional Neural Networks." In MLIS@ICML, pp. 21-25. 2015.

? Self-supervised

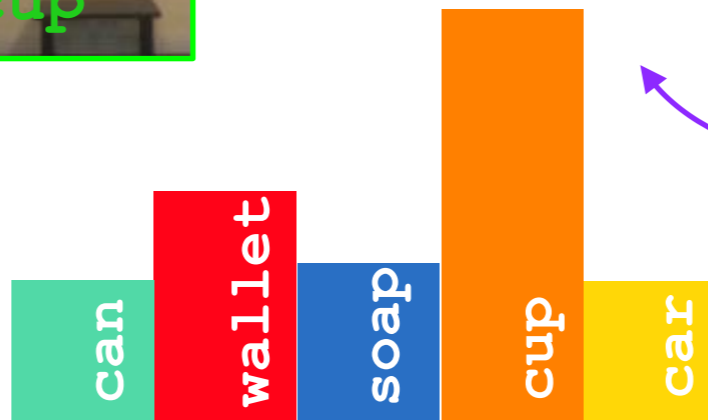

? Reliable

? Exploits
contextual
information

? Learns
incrementally



© Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.



iCubWorld28 Dataset

Overview

2014: "Household"



laundry
detergent

plate

dishwashing
detergent

sponge

cup

soap

sprayer

7 categories
4 objects/category
→ 28 objects

4 acquisitions

Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale. Used with permission.

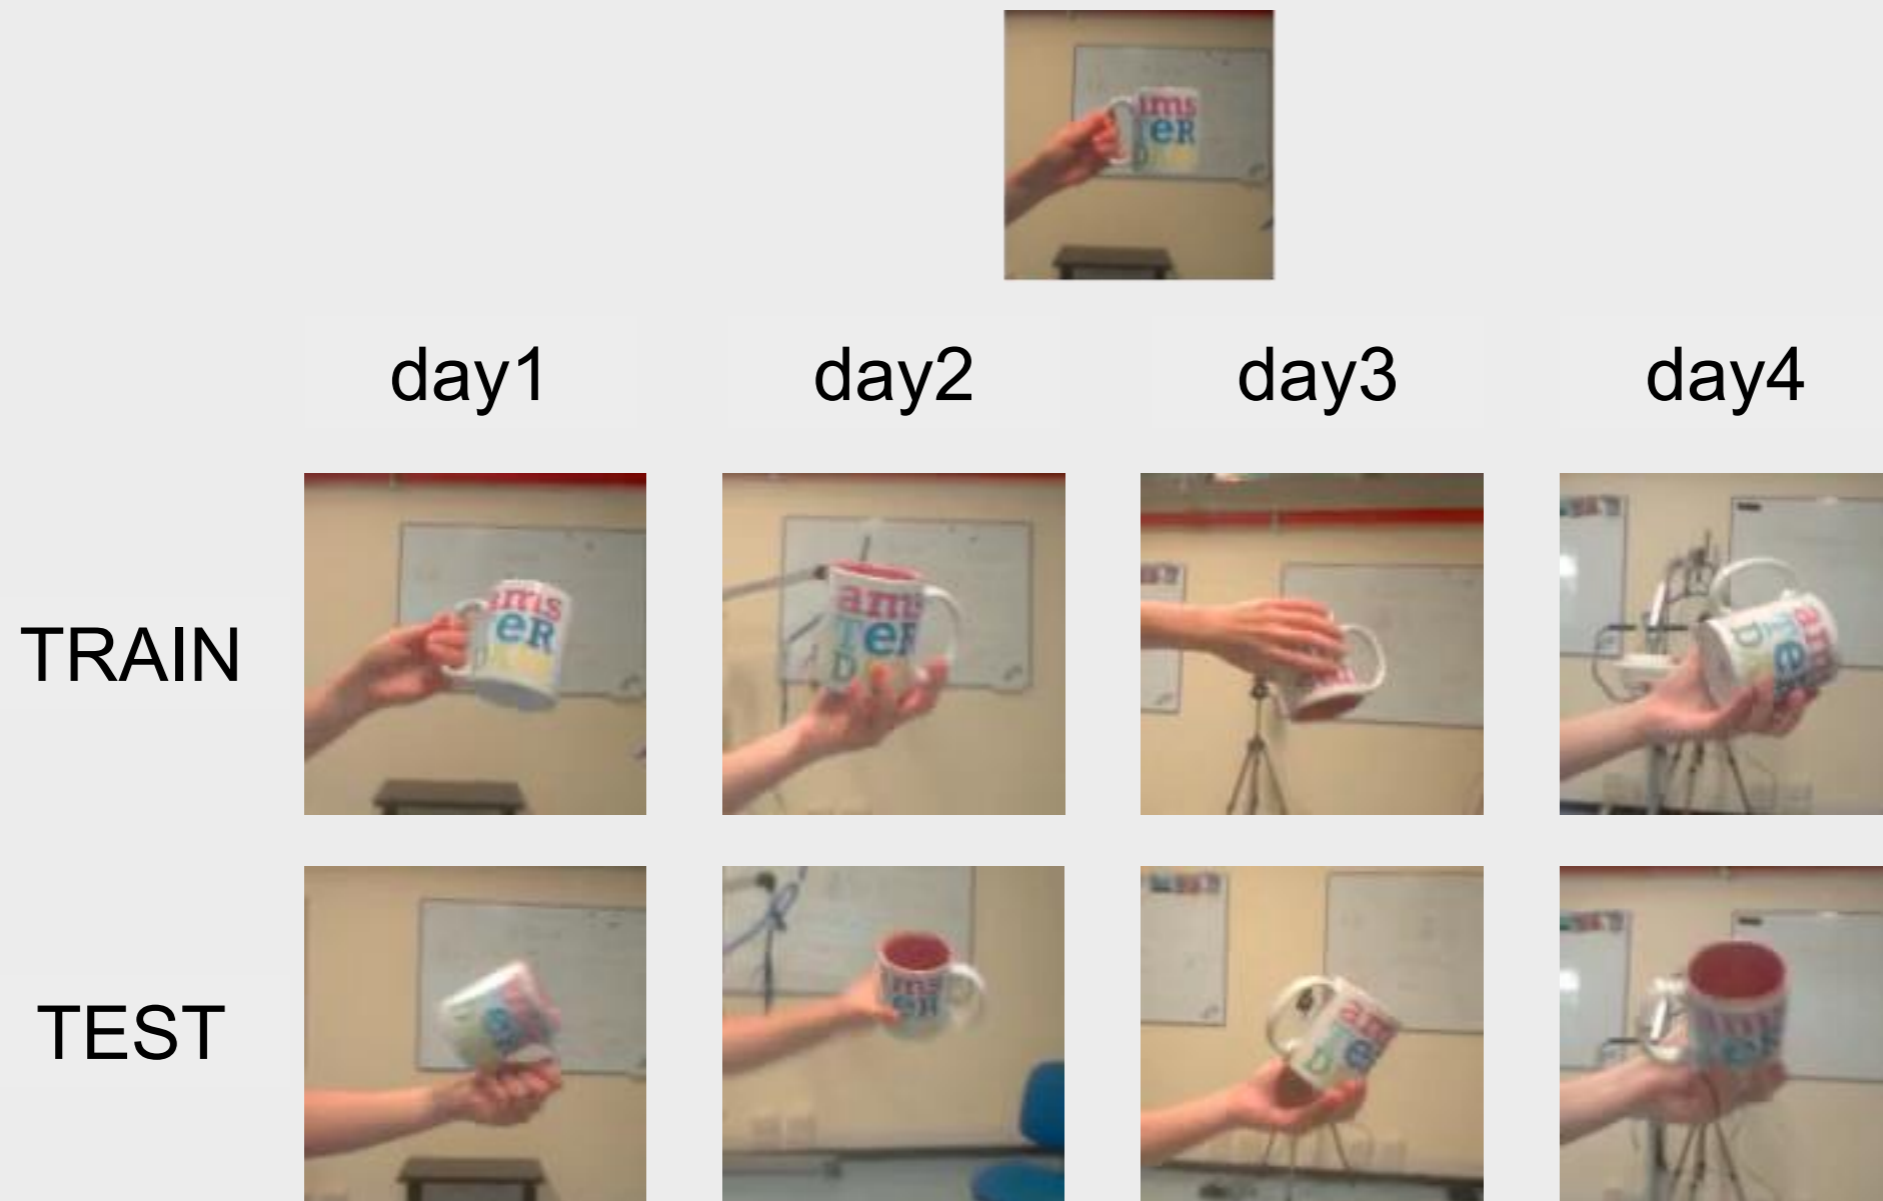
Benchmarking deep Conv Nets for Real-world Object Recognition: How many Objects can iCub Learn?

[arXiv: 1504.03154](https://arxiv.org/abs/1504.03154),

iCubWorld28 Dataset

Examples of Acquired Videos

2014: “Household”



© Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale.
All rights reserved. This content is excluded from our Creative Commons license. For more
information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Benchmarking deep Conv Nets for Real-world Object Recognition: How many Objects can iCub Learn?
arXiv: 1504.03154,

iCubWorld28 Dataset

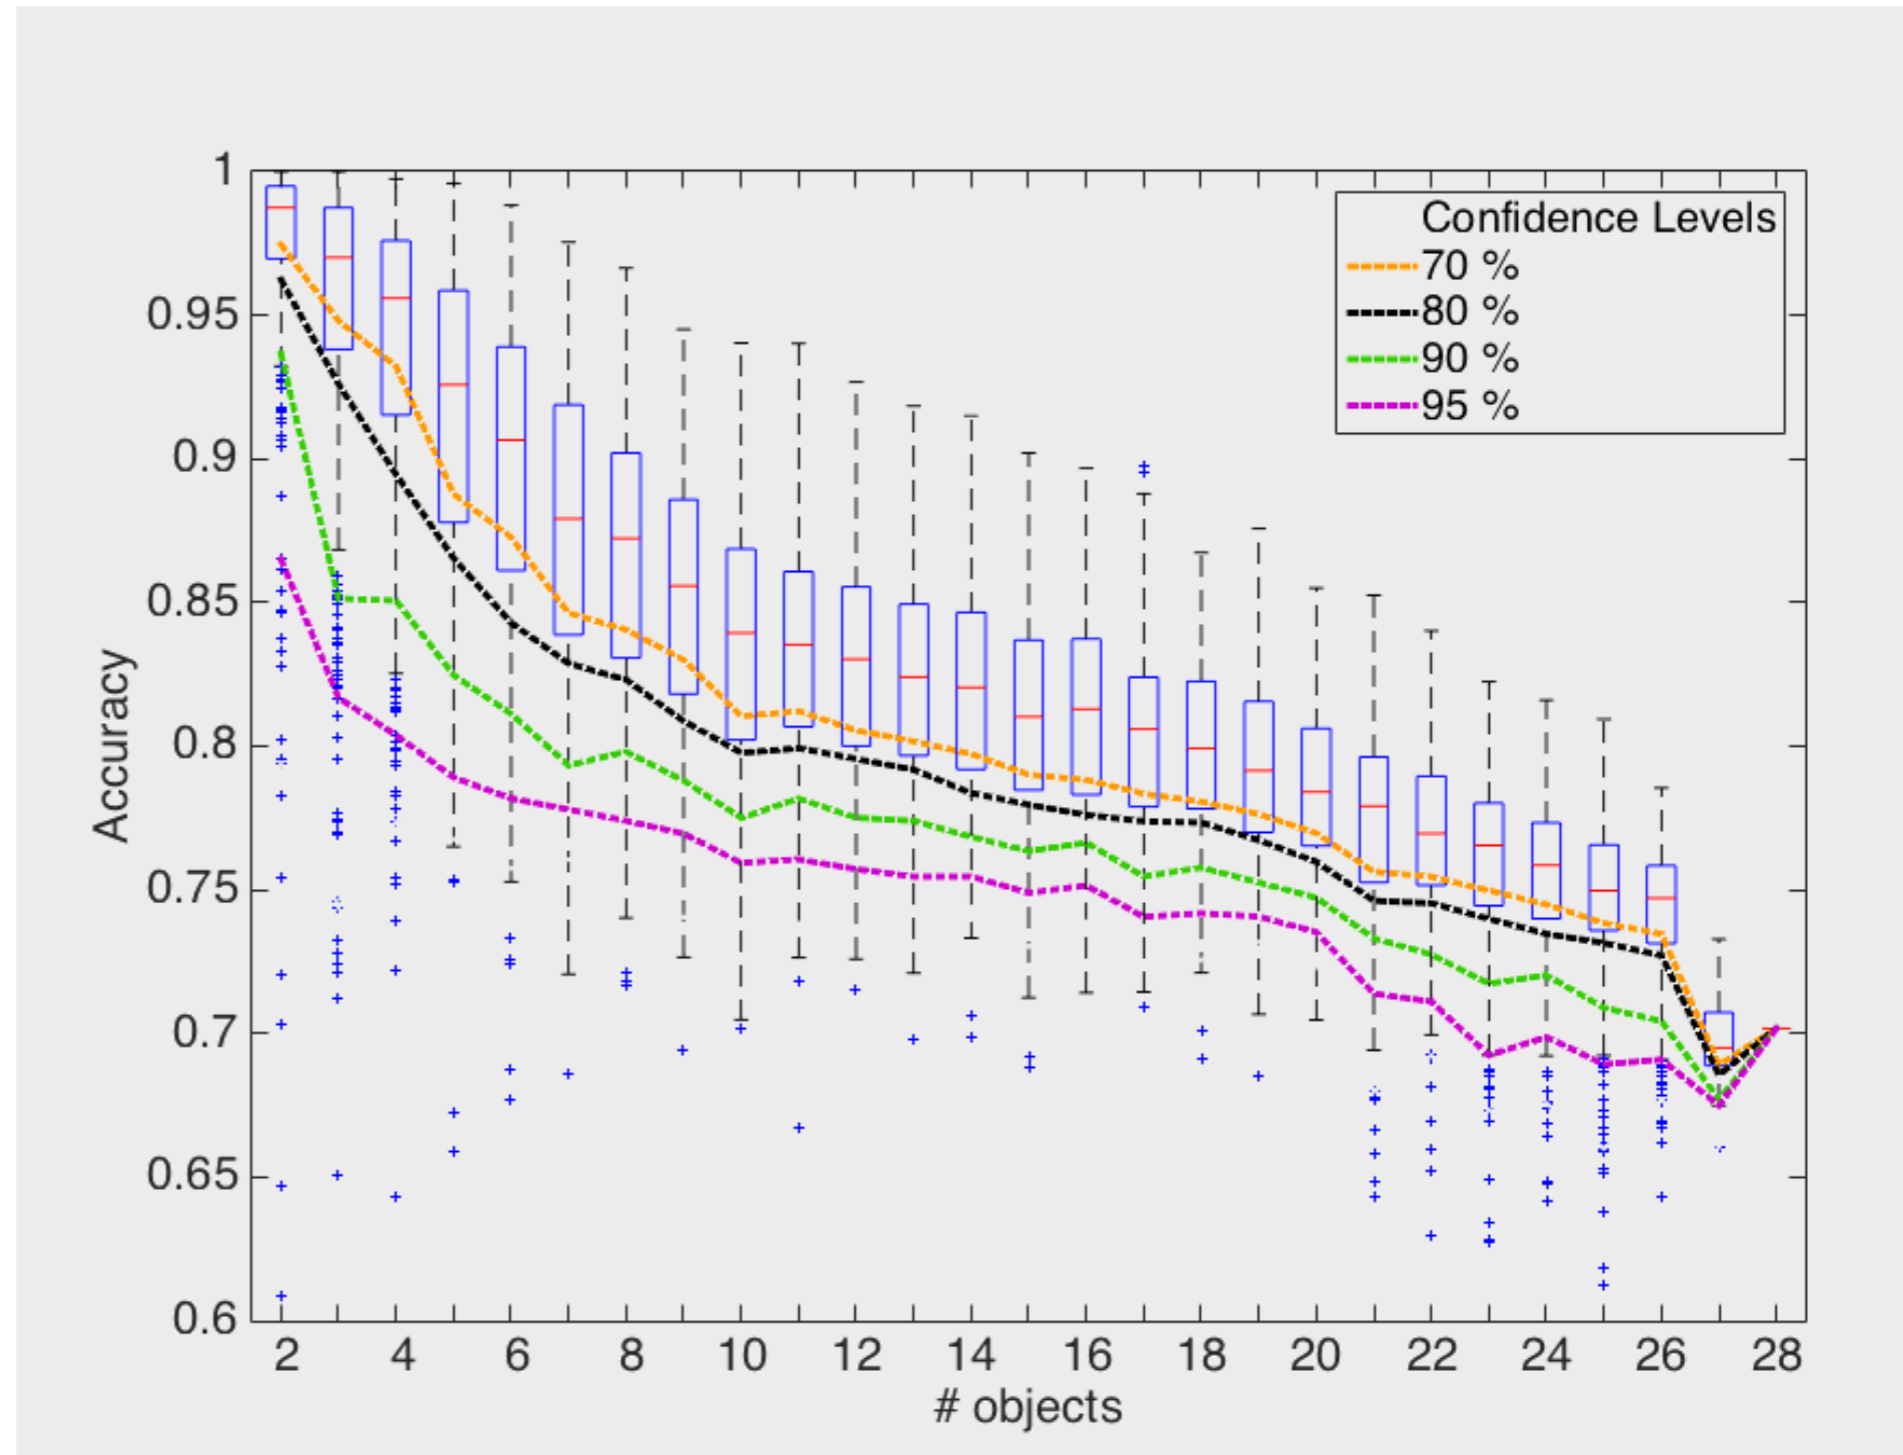
Object Identification “Data Sheet”

? Self-supervised

? Reliable

? Exploits contextual information

? Learns incrementally



Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale. Used with permission.

Benchmarking deep Conv Nets for Real-world Object Recognition: How many Objects can iCub Learn?

[arXiv: 1504.03154](https://arxiv.org/abs/1504.03154),

iCubWorld28 Dataset

Clutter and Scale

? Self-supervised

? Reliable

? Exploits contextual information

? Learns incrementally



Image



Crop 1



Crop 2



Manual

		TEST Accuracy (%)			
		Image	Crop1	Crop2	Manual
TRAIN	Image	50.6	48.8	36.3	20.6
	Crop1	50.3	62.2	57.7	24.9
	Crop2	30.1	50.8	73.9	28.7
	Manual	6.8	8.9	12.2	81.7

Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale. Used with permission.

Benchmarking deep Conv Nets for Real-world Object Recognition: How many Objects can iCub Learn?

[arXiv: 1504.03154](https://arxiv.org/abs/1504.03154),

iCubWorld28 Dataset

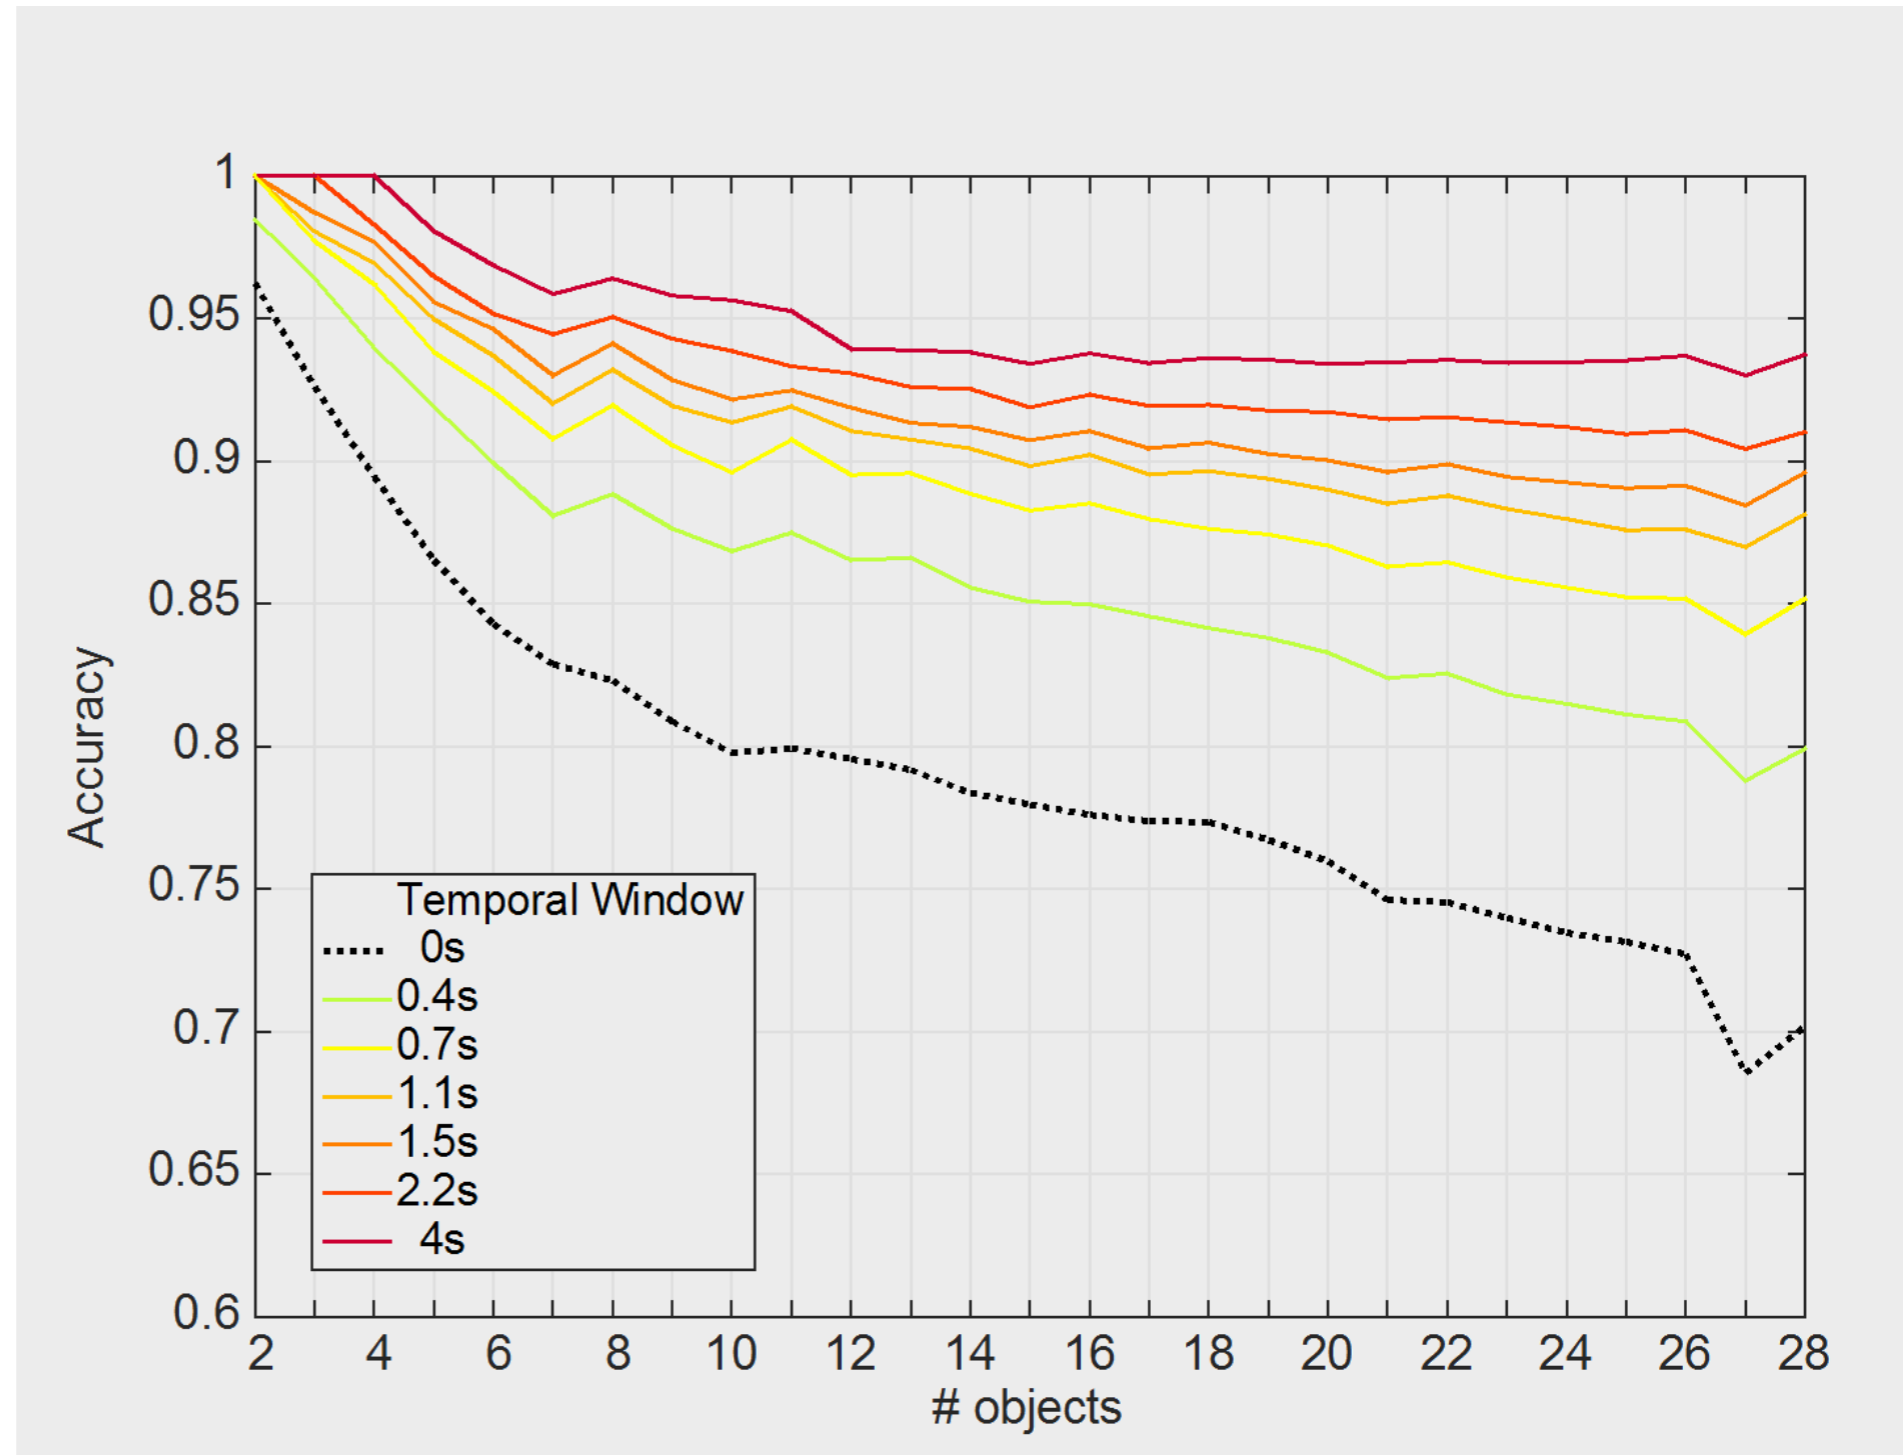
Temporal Contextual Information

? Self-supervised

? Reliable

? Exploits contextual information

? Learns incrementally



Courtesy of Giulia Pasquale, Carlo Ciliberto, Francesca Odone, Lorenzo Rosasco and Lorenzo Natale. Used with permission.

Benchmarking deep Conv Nets for Real-world Object Recognition: How many Objects can iCub Learn?

[arXiv: 1504.03154](https://arxiv.org/abs/1504.03154),

iCubWorld Datasets

Ongoing Work

2015:

“Kitchen” + “Food” + “Toys” + “Bathroom” + “Daily use” + “Office” + “Tools”

oven glove	bottle	beach:	body cream	sunglasses	mouse	paint brush
squeezer	box	shovel	hair brush	hair clip	organizer	scissors
cup	can	bucket	soap	watch	calculator	scotch
		toy rake				

1. Object Categorization Dataset
2. Continuously Expandable in Time

21 categories
10 objects/category
→ 200 objects

3. Tagged by Nuisance Factors

5 acquisitions
divided by nuisance:

- ✓ scale
- ✓ 2D rotation
- ✓ 3D rotation
- ✓ translation
- ✓ mixed

4. Depth information available (left+right cameras)

iCubWorld Datasets

Disparity-driven segmentation



Courtesy of Giulia Pasquale, Tanis Mar, Carlo Ciliberto, Lorenzo Rosasco, and Lorenzo Natale. Used with permission.

Enabling Depth-driven Visual Attention on the iCub robot: Instructions for Use and New Perspectives
submitted to Humanoids 2015

iCubWorld Datasets

Ongoing Work

2015:

“Kitchen” + “Food” + “Toys” + “Bathroom” + “Daily use” + “Office” + “Tools”

oven glove
squeezer
cup

bottle
box
can

shovel
bucket
toy rake

body cream
hair brush
soap

sunglasses
hair clip
watch

mouse
organizer
calculator

paint brush
scissors
scotch



iCubWorld Datasets

Ongoing Work

2015:

“Kitchen” + “Food” + “Toys” + “Bathroom” + “Daily use” + “Office” + “Tools”


translation



scale



mixed



2D rotation



3D rotation



**Application & Data
are available for
projects 5.2 & 5.3!!**



MIT OpenCourseWare
<https://ocw.mit.edu>

Resource: Brains, Minds and Machines Summer Course
Tomaso Poggio and Gabriel Kreiman

The following may not correspond to a particular course on MIT OpenCourseWare, but has been provided by the author as an individual learning resource.

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.