Neuro 140:
Biological & Artificial Intelligence

Feb 9th, 2021: **Using AI in the laboratory**

Mackenzie Mathis, PhD

Bertarelli Foundation Chair of Integrative Neuroscience

Assistant Professor, EPFL

mackenzie@post.harvard.edu

European Laboratory for Learning and Intelligent Systems

Outline

3 – 5 PM

(1) Animal behavior in the laboratory: what problems could use “better” quantification of behavior?

(2) What are the challenges for video based analysis?

(3) How can deep neural networks help automate the process

(4) Transfer learning & pose estimation

~4 – 4:15 PM BREAK!

(5) Using transfer learning for robustness in o.o.d. data

(6) Google Colab intro / model zoo (if you have a dog/cat video, you can have it handy!)

E. Muybridge, 1887 (zoopraxiscope)

Ota et al. 2015 Sci Reports

Mathis et al 2017 Neuron

Animal behavior in the lab: how to quantify the behavior?

Natural behaviors

Diverse species, natural environments, many animals

Looming stimulus

M. Meister/Cal Tech

Multi-directional reaching
(and neural reconstruction with DL)

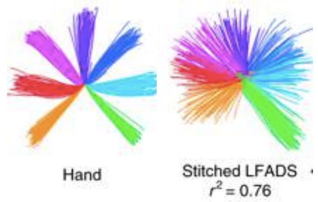

LFADS – Pandarinath et al 2018

No trial structure

Classical conditioning

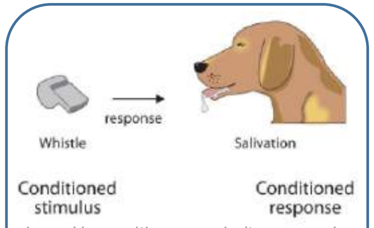

http://open.lib.umn.edu/intropsyc/

Robust & “simple” trial-based behaviors

Complex & more variable (trial-based) behaviors

Measuring behavior

observation → **pose estimation** → kinematics

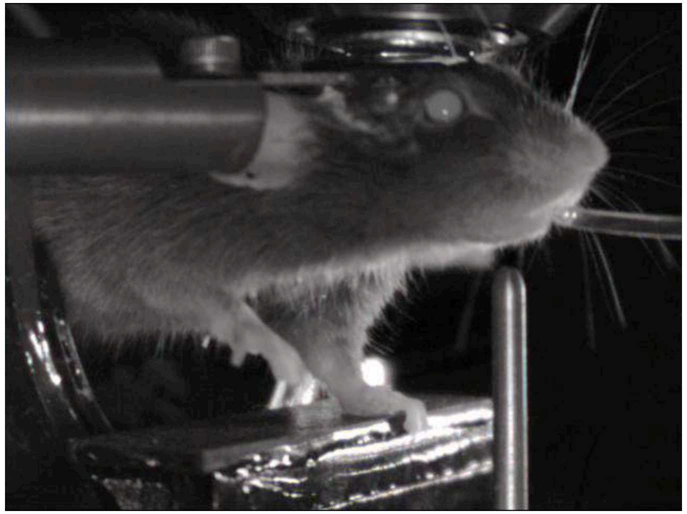

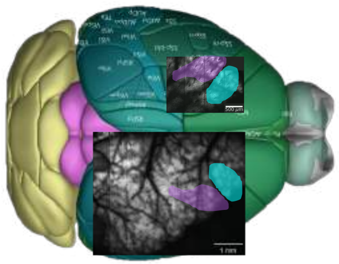

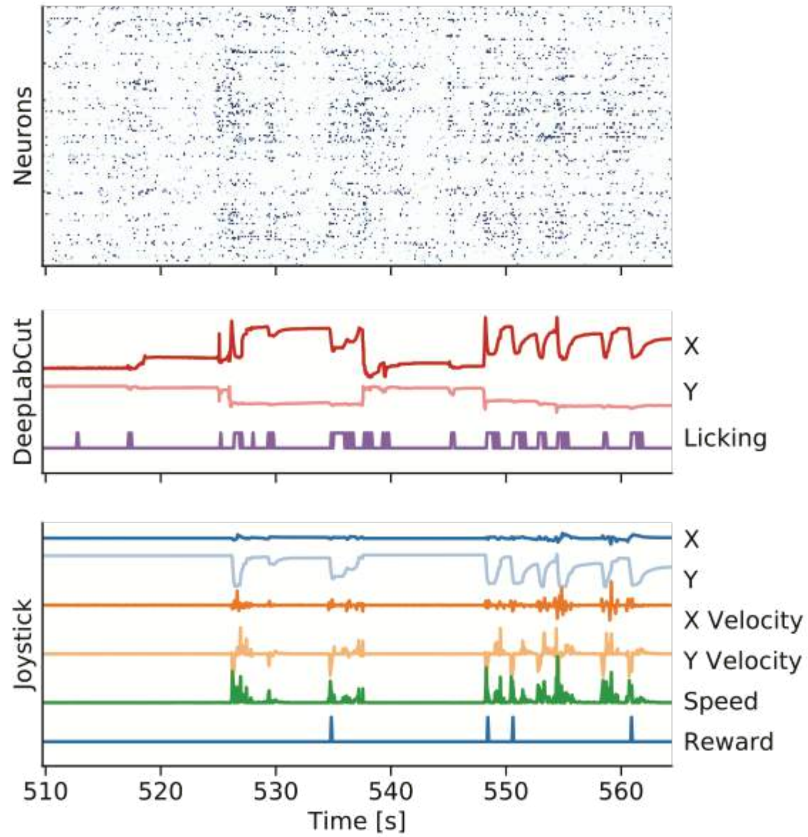

Pose estimation in the laboratory

Fransechi et al. 2016
Wilschko et al. 2015
Azim et al. 2014
Kawai et al. 2015
Dell et al. 2014
Berman 2018
Vargas-Irwin et al. 2010
Pereira et al, 2018
Human motion capture systems

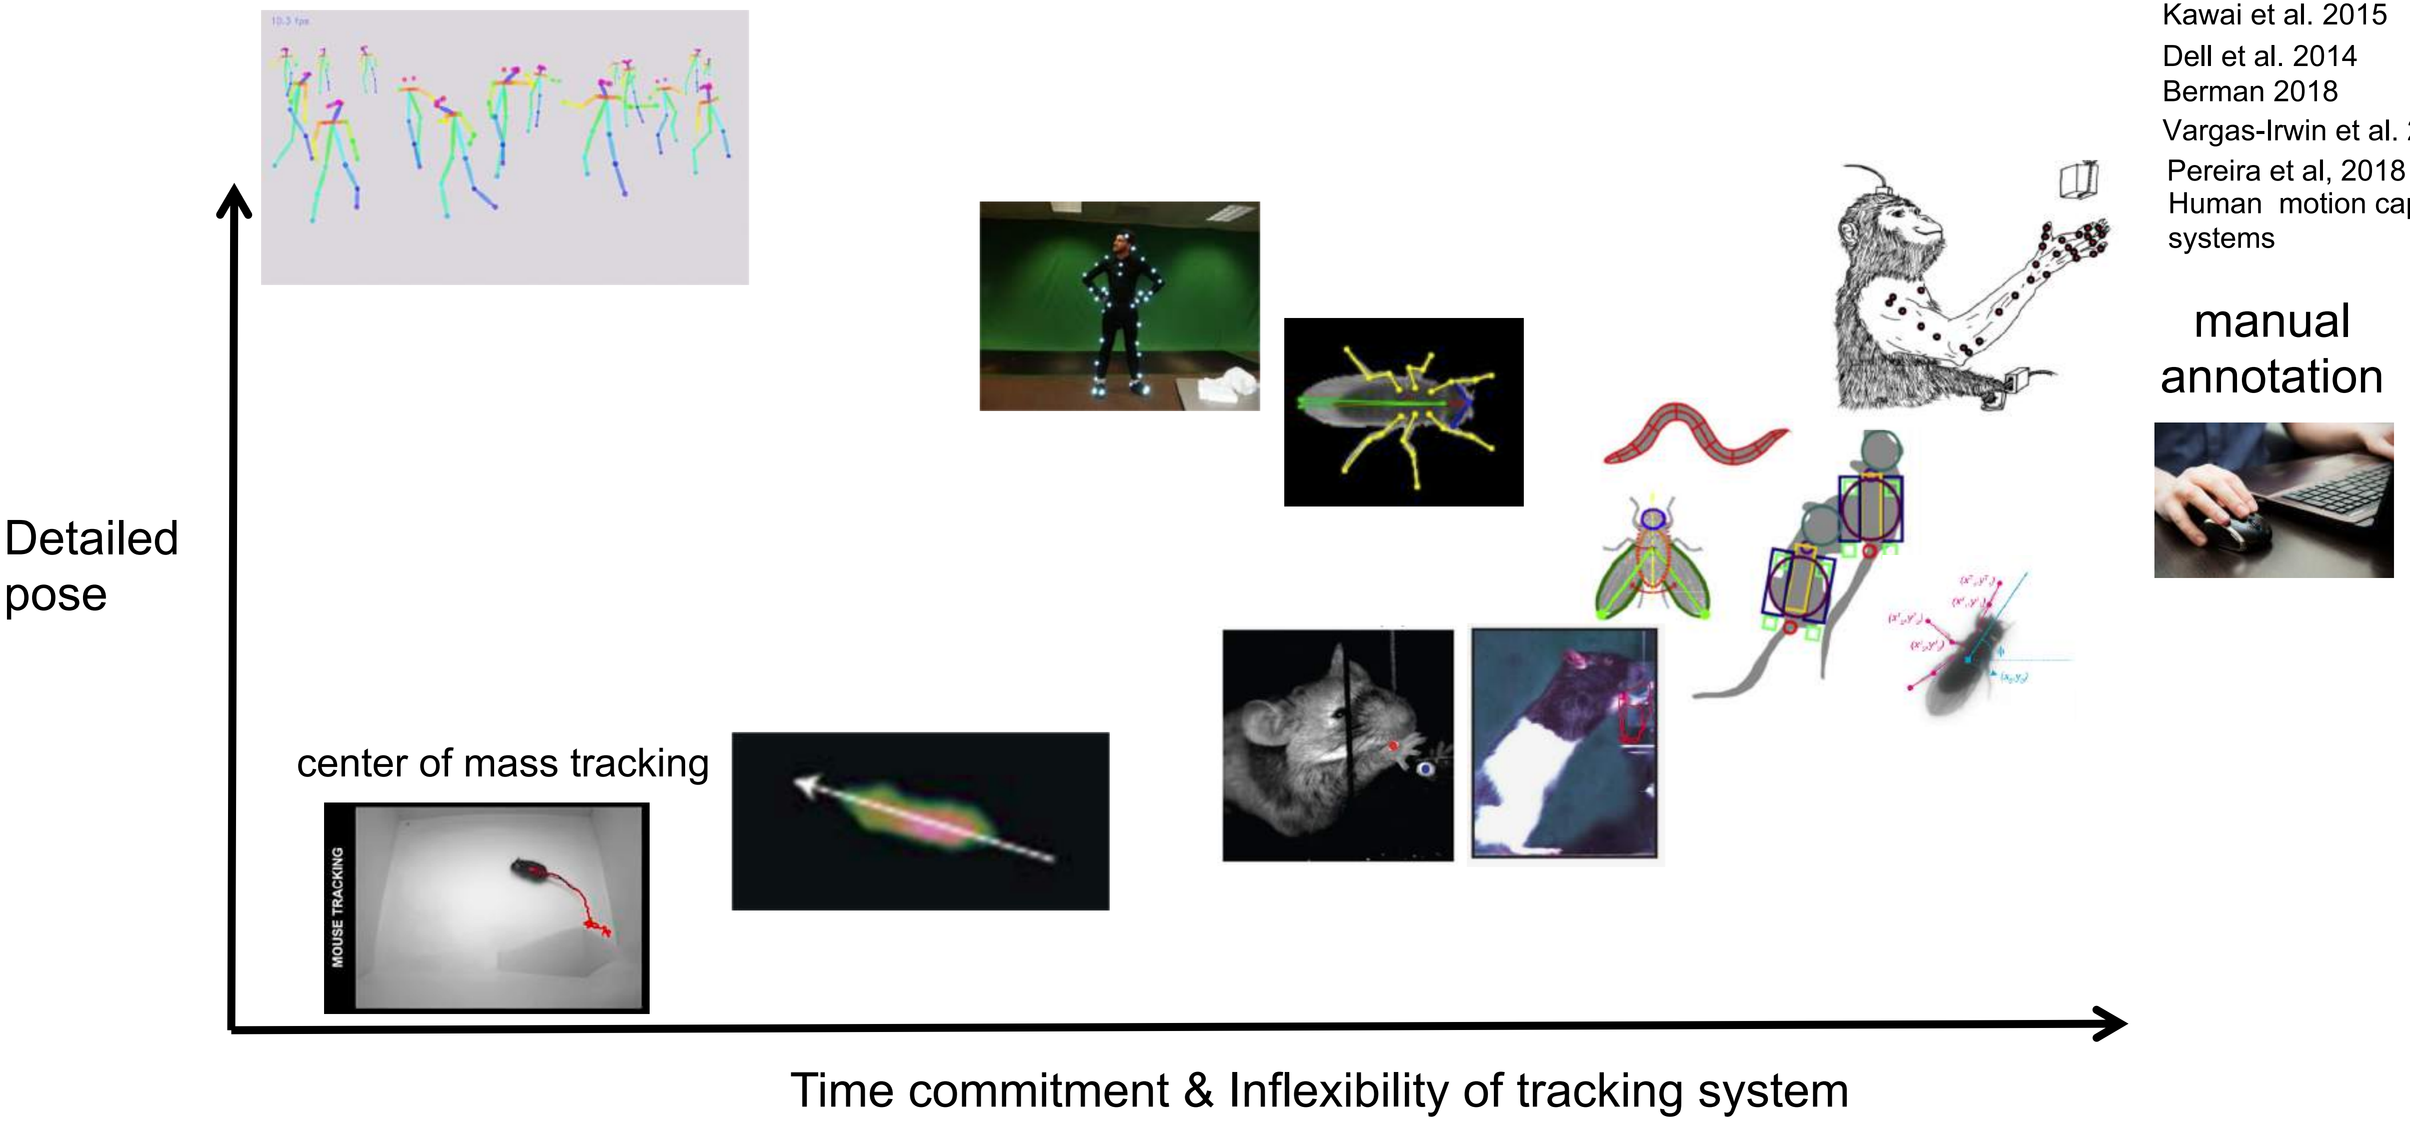

Deep learning in the lab

DeepPose
DeeperCut
OpenPose
Conv. PoseMachines
…

deep neural networks

image ⟶ Predictor ⟶ pose

train

A lot of labeled images (>10^6 joints!)

Andrew Ng

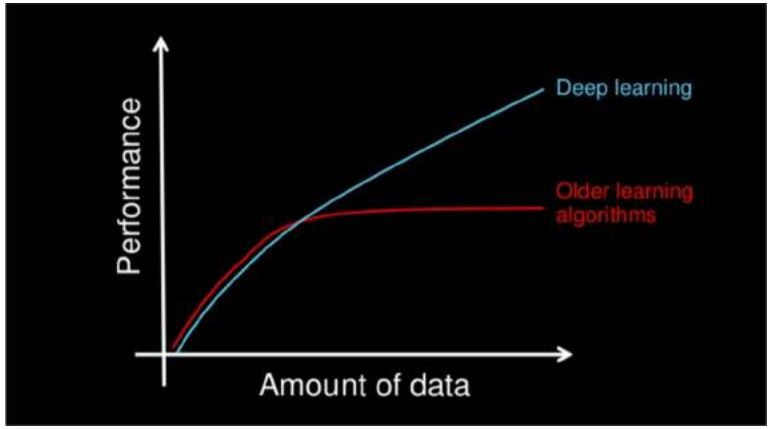

DATA hungry algorithms… how to bring this to the lab?

Transfer Learning: take a trained network and ask it to learn a new task

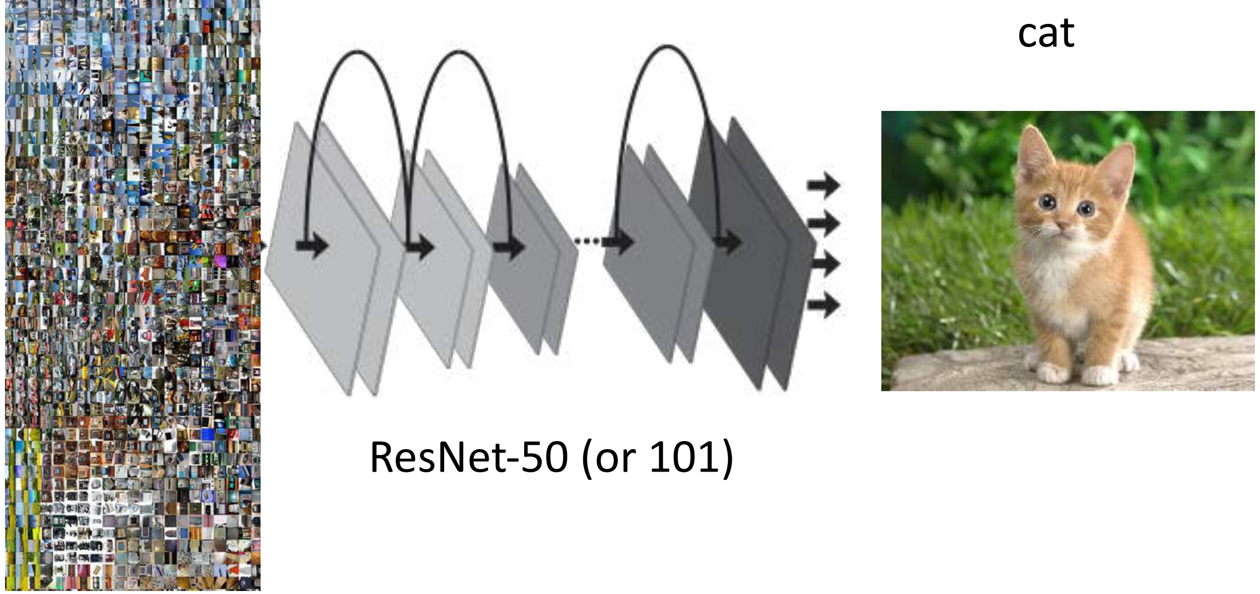

ImageNet

Flute

Strawberry

Traffic light

Backpack

Matchstick

Bathing cap

Sea lion

Racket

IMAGENET

Olga Russakovsky*, Jia Deng*, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg and Li Fei-Fei. (* = equal contribution) **ImageNet Large Scale Visual Recognition Challenge**. *International Journal of Computer Vision*, 2015.

CNNs

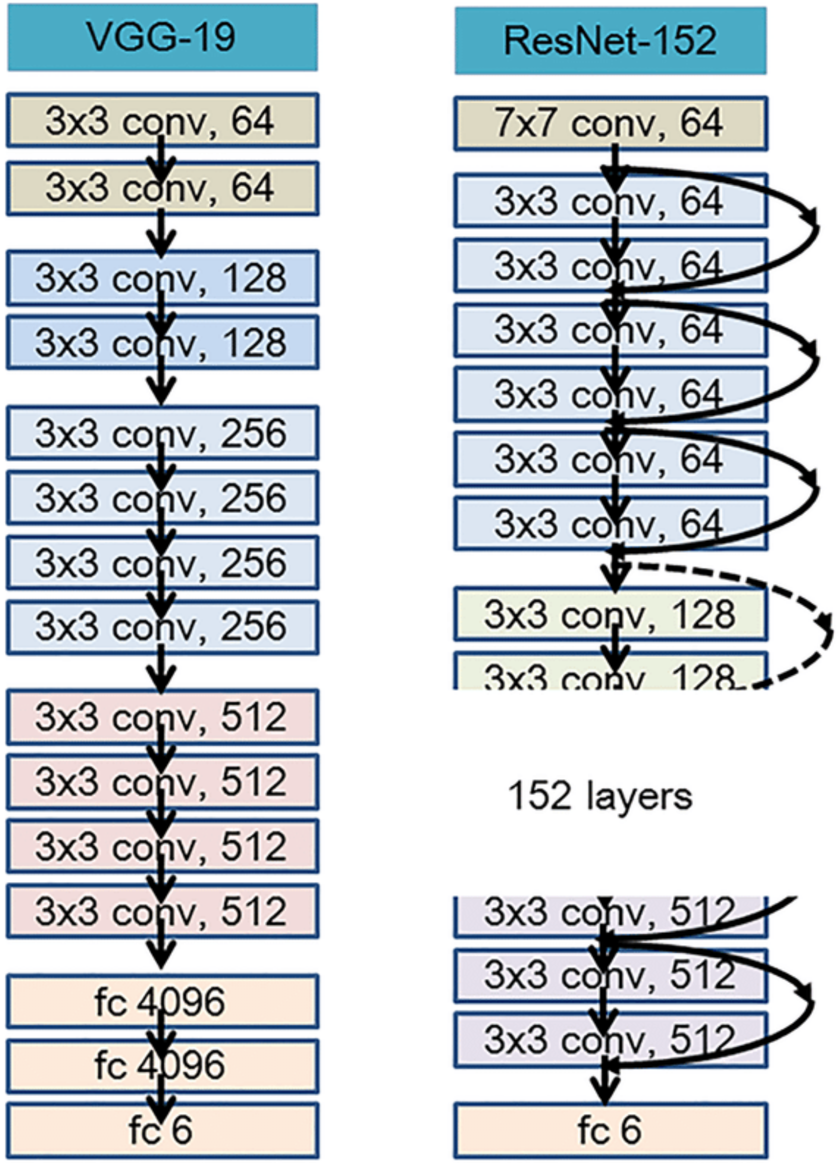
VGG-19
3x3 conv, 64
3x3 conv, 64
3x3 conv, 128
3x3 conv, 128
3x3 conv, 256
3x3 conv, 256
3x3 conv, 256
3x3 conv, 256
3x3 conv, 512
3x3 conv, 512
3x3 conv, 512
3x3 conv, 512
fc 4096
fc 4096
fc 6
ResNet-152
7x7 conv, 64
3x3 conv, 64
3x3 conv, 64
3x3 conv, 64
3x3 conv, 64
3x3 conv, 64
3x3 conv, 64
3x3 conv, 128
152 layers
3x3 conv, 512
3x3 conv, 512
3x3 conv, 512
fc 6

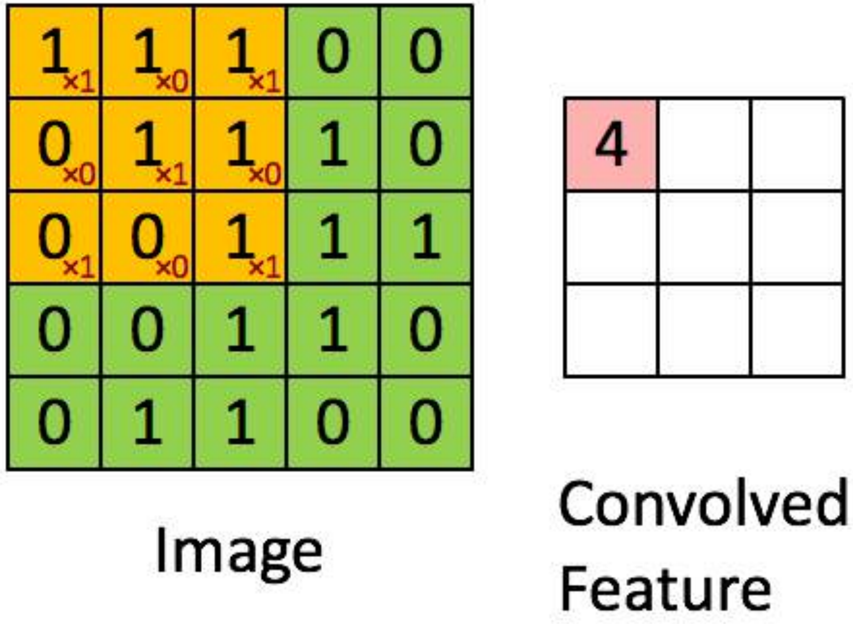
1 ×1 1 ×0 1 ×1 0 0
0 ×0 1 ×1 1 ×0 1 0
0 ×1 0 ×0 1 ×1 1 1
0 0 1 1 0
0 1 1 0 0
Image
4
Convolved Feature

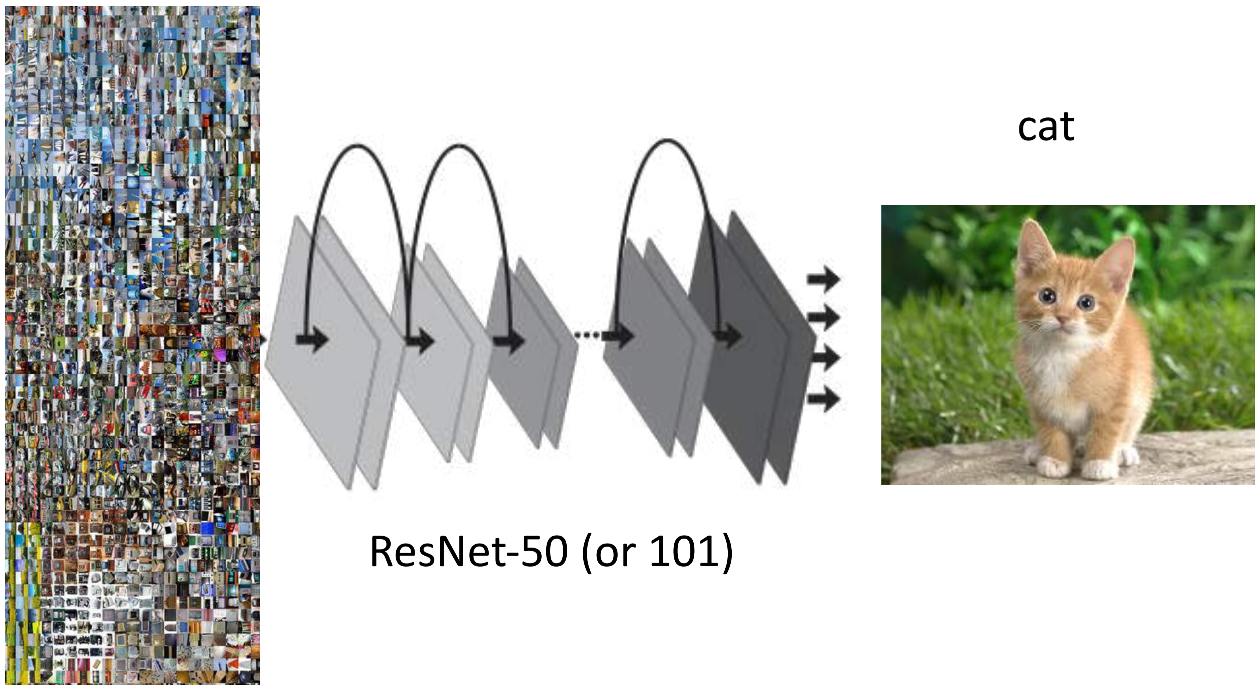

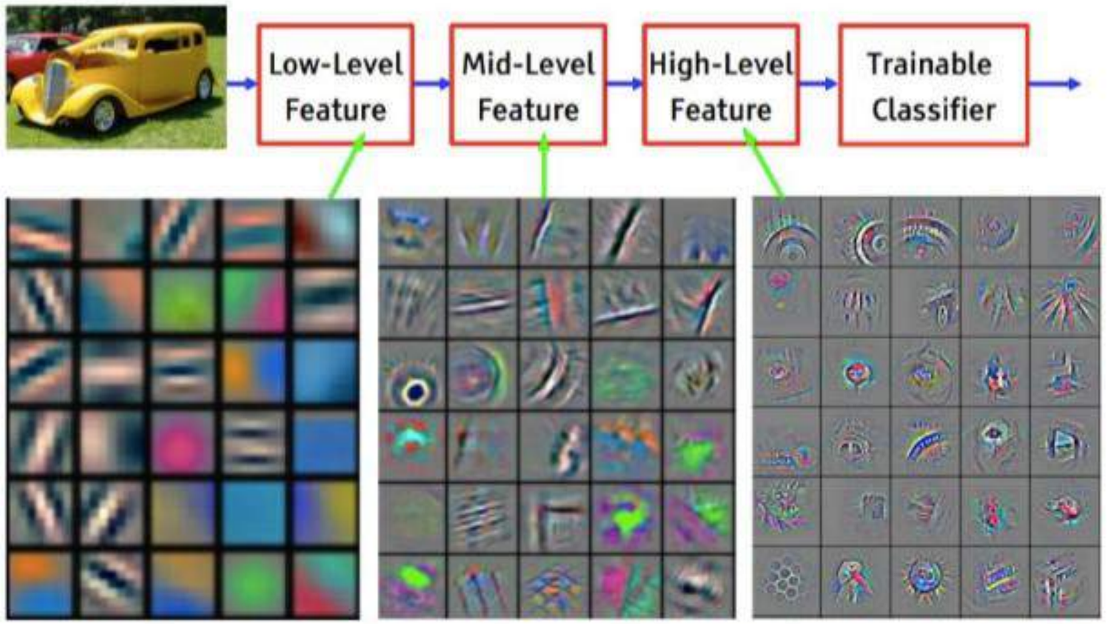

Feature visualization of convolutional net trained on ImageNet from [Zeiler & Fergus 2013]

Do Better ImageNet Models Transfer Better?

Simon Kornblith*, Jonathon Shlens, and Quoc V. Le
Google Brain
{skornblith,shlens,qvl}@google.com

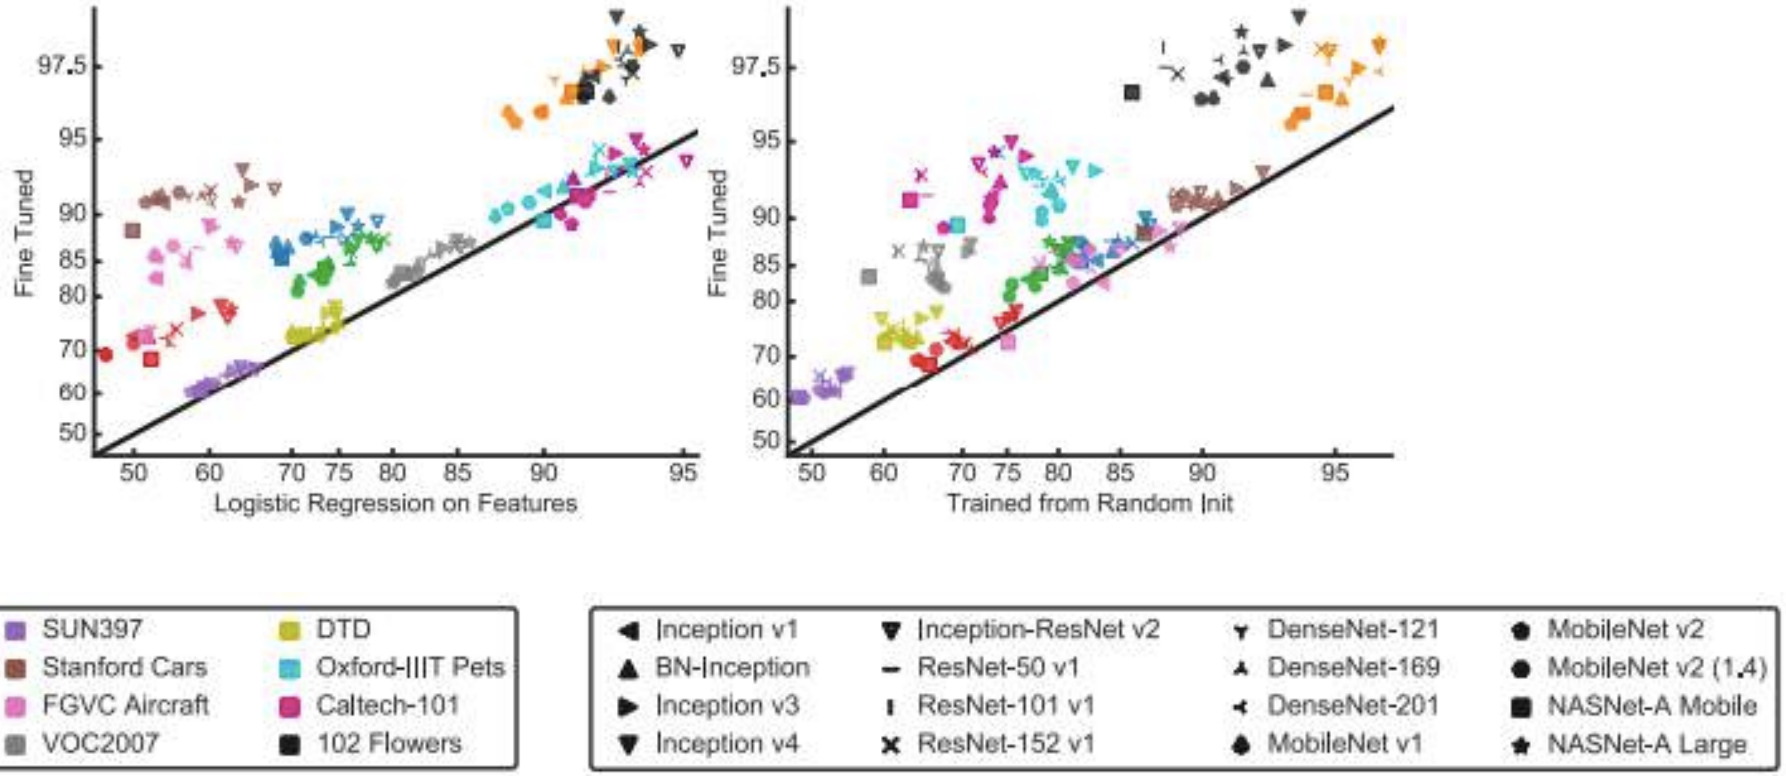

Do Better ImageNet Models Transfer Better?

Simon Kornblith*, Jonathon Shlens, and Quoc V. Le
Google Brain
{skornblith,shlens,qvl}@google.com

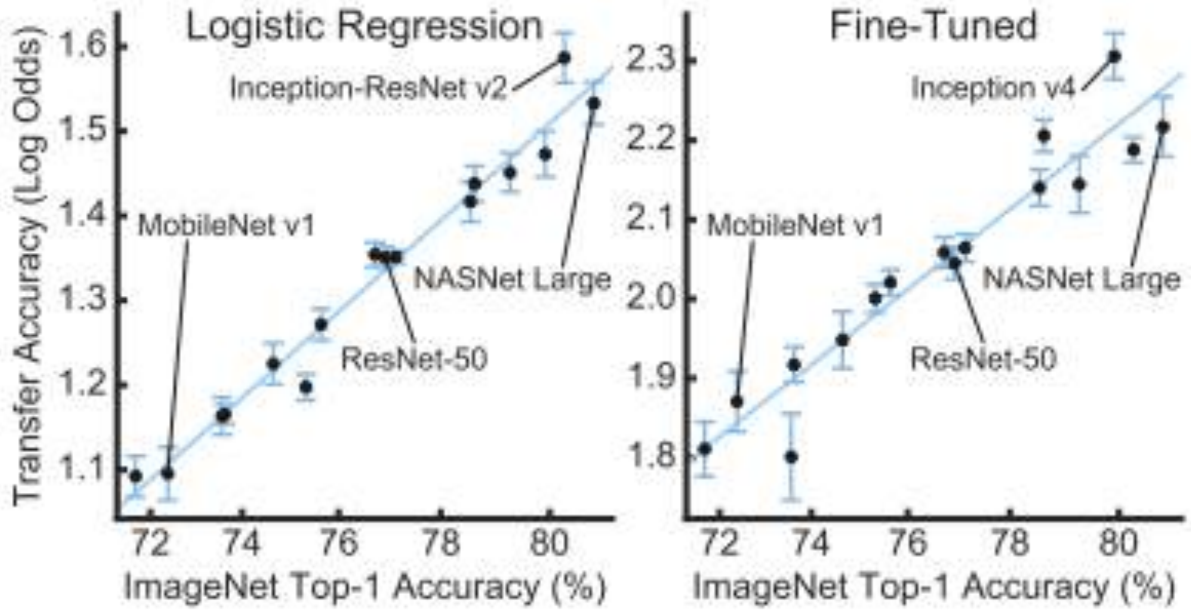

Figure 1. Transfer learning performance is highly correlated with ImageNet top-1 accuracy for fixed ImageNet features (left) and fine-tuning from ImageNet initialization (right). The 16 points in each plot represent transfer accuracy for 16 distinct CNN architectures, averaged across 12 datasets after logit transformation (see Section 3). Error bars measure variation in transfer accuracy across datasets. These plots are replicated in Figure 2 (right).

Steel drum

Output:
Scale
T-shirt
Steel drum
Drumstick
Mud turtle

✓

Output:
Scale
T-shirt
Giant panda
Drumstick
Mud turtle

✗

DEEP GAZE I: BOOSTING SALIENCY PREDICTION WITH FEATURE MAPS TRAINED ON IMAGENET

Matthias Kümmerer, Lucas Theis & Matthias Bethge
Werner Reichardt Centre for Integrative Neuroscience
University Tübingen, Germany
{matthias.kuemmerer,lucas,matthias}@bethgelab.org

DeepGaze II: Reading fixations from deep features trained on object recognition

Matthias Kümmerer Thomas S. A. Wallis
Matthias Bethge

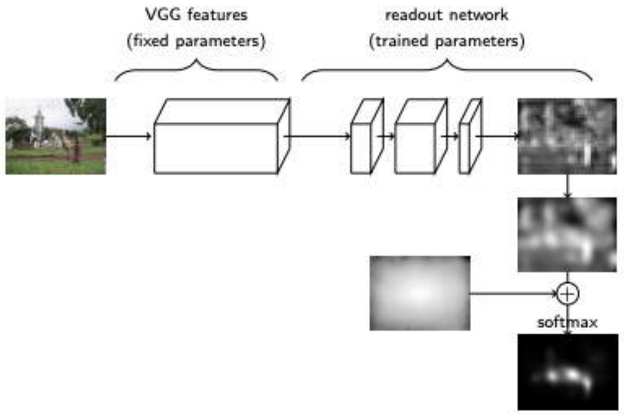

Figure 1: The architecture of DeepGaze II. The activations of a subset of the VGG feature maps for a given image are passed to a second neural network (the readout network) consisting of four layers of 1×1 convolutions. The parameters of VGG are held fixed through training (only the readout network learns about saliency prediction). This results in a final saliency map, which is then blurred, combined with a centre bias and converted into a probability distribution by means of a softmax.

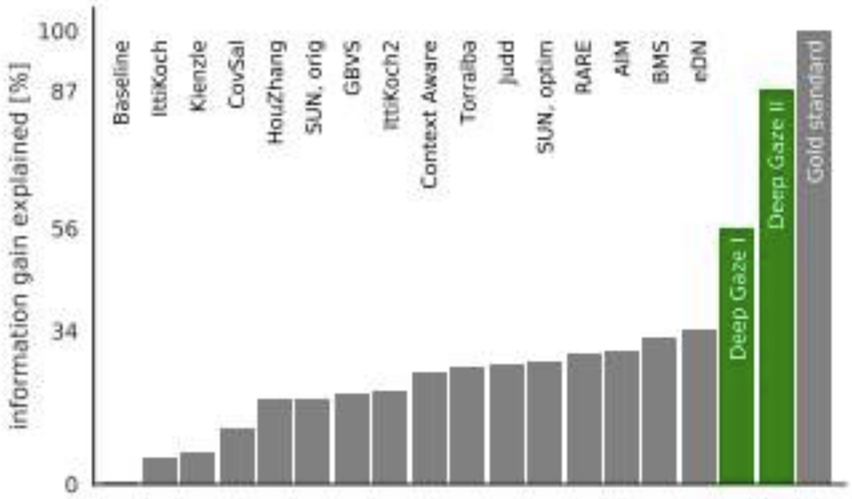

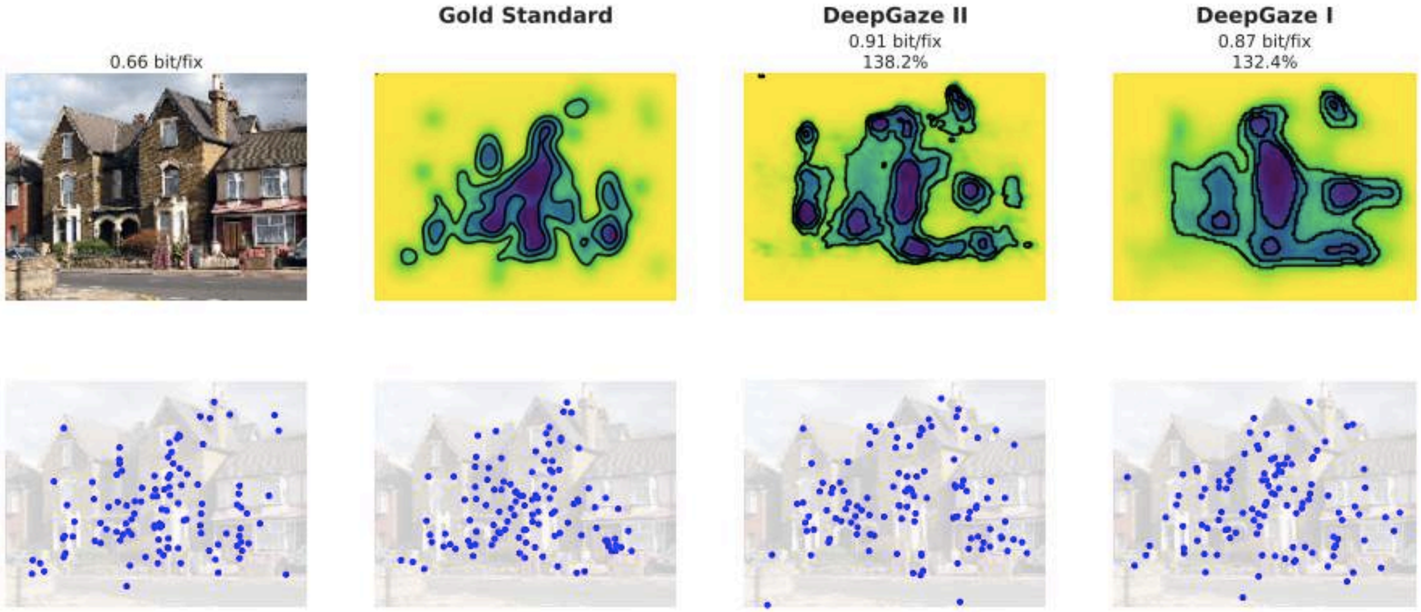

Challenges for animal pose estimation in the laboratory

- animals have highly different bodies
- not practical to label >10,000 frames for training
- 3D postures
- fast video analysis for many experiments (such as closed-loop optogenetics)
- Multi-animal tracking
- Generalization & Robustness

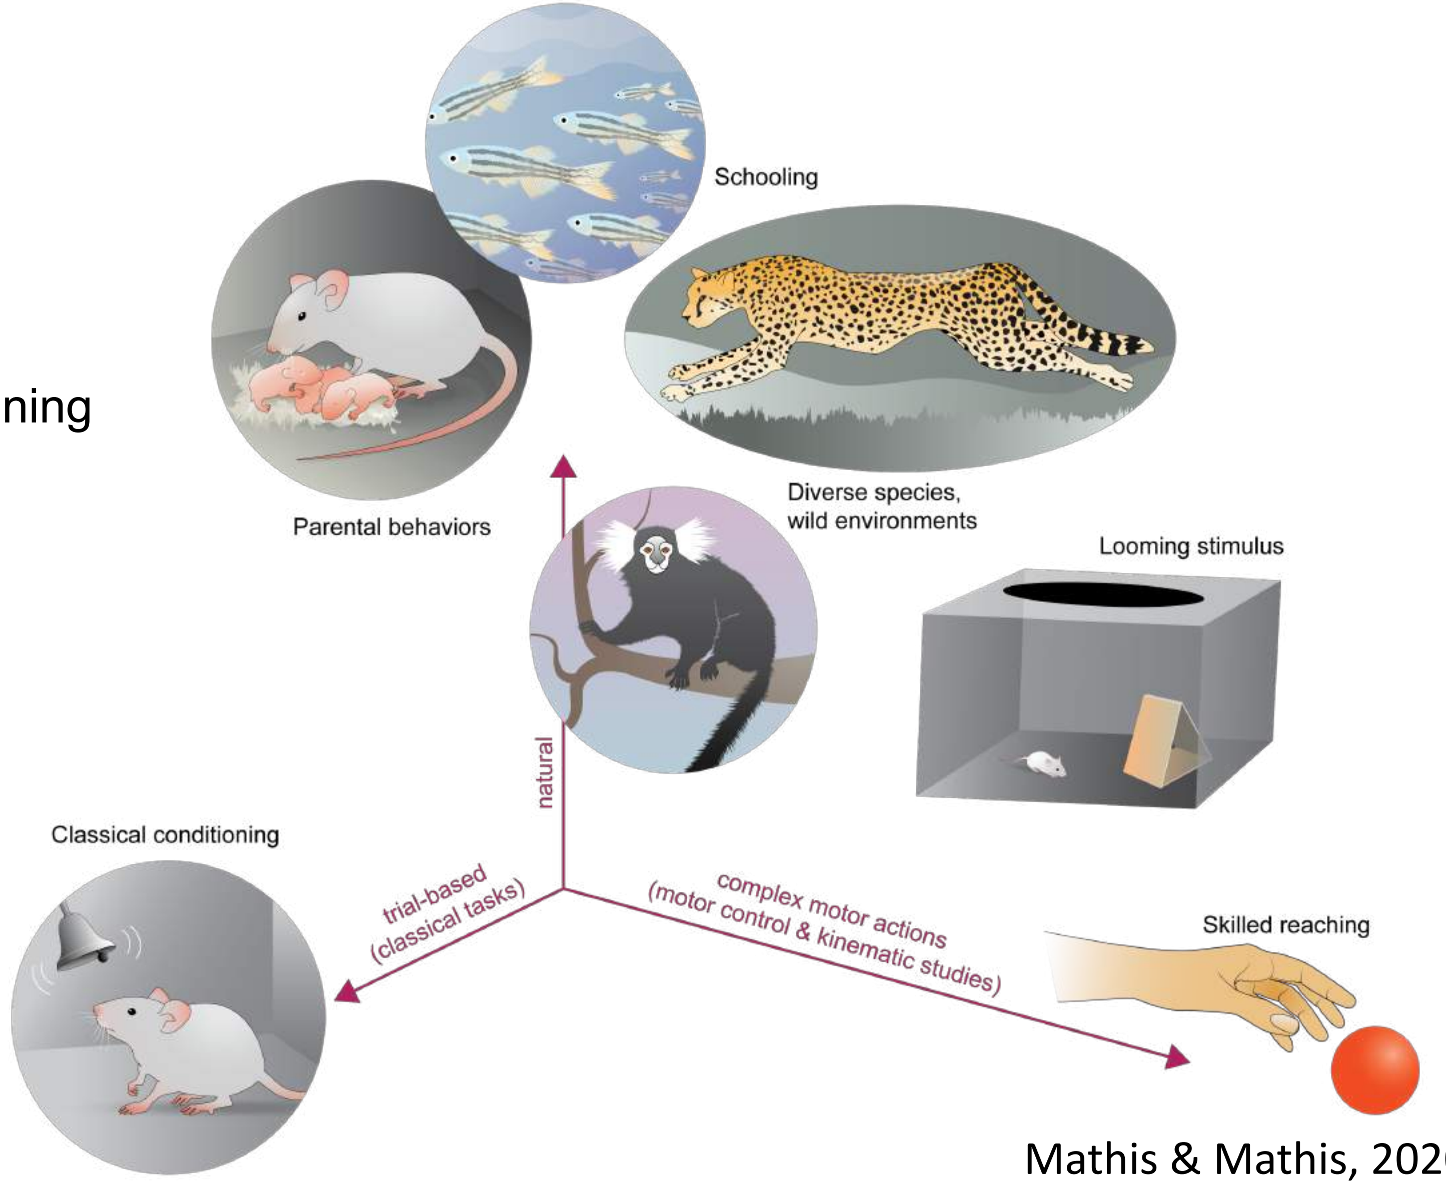

Mathis & Mathis, 2020
Current Opinion in Neurobiology

DeepLabCut: animal pose estimation w/transfer learning

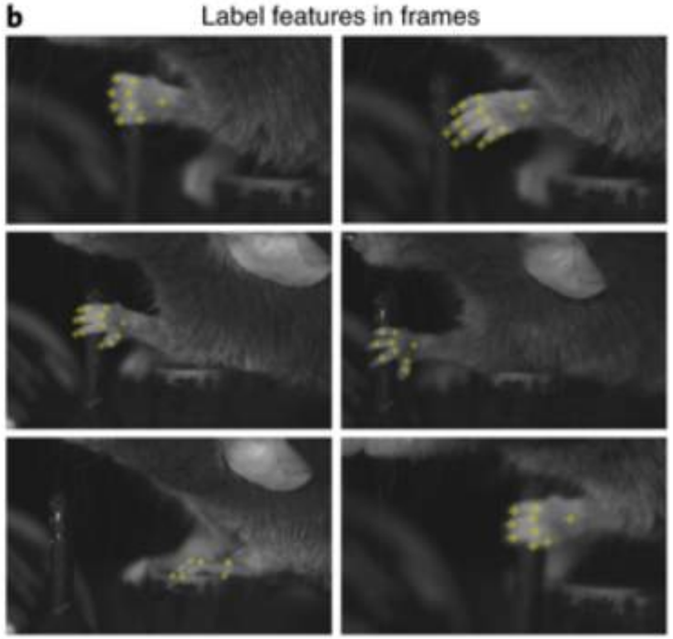

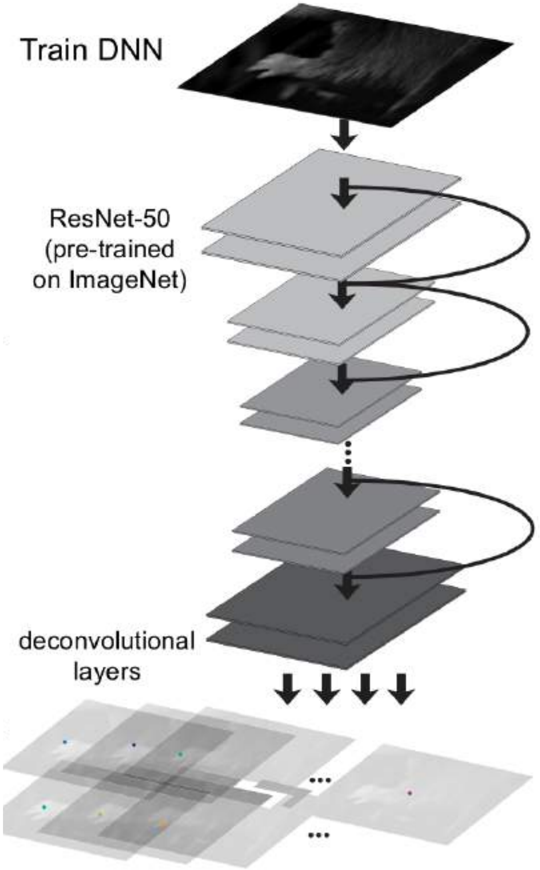

A. Mathis …. M.W. Mathis* and Bethge*
Nature Neurosci, 2018

DeepLabCut: a toolbox for markerless pose estimation

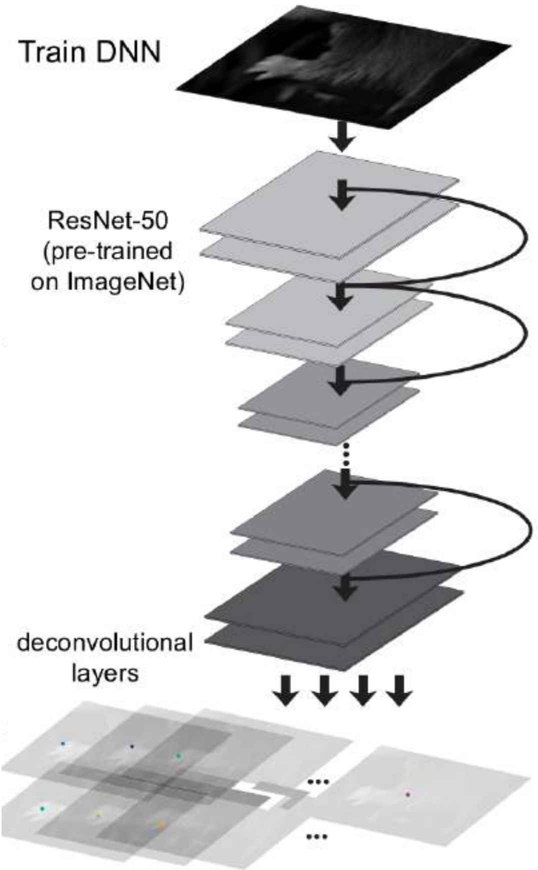

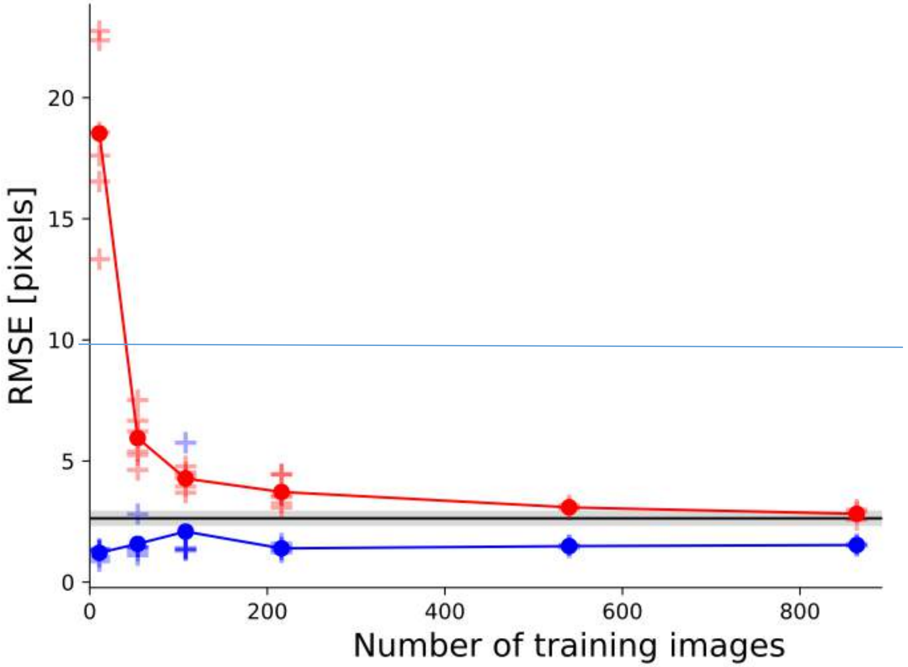

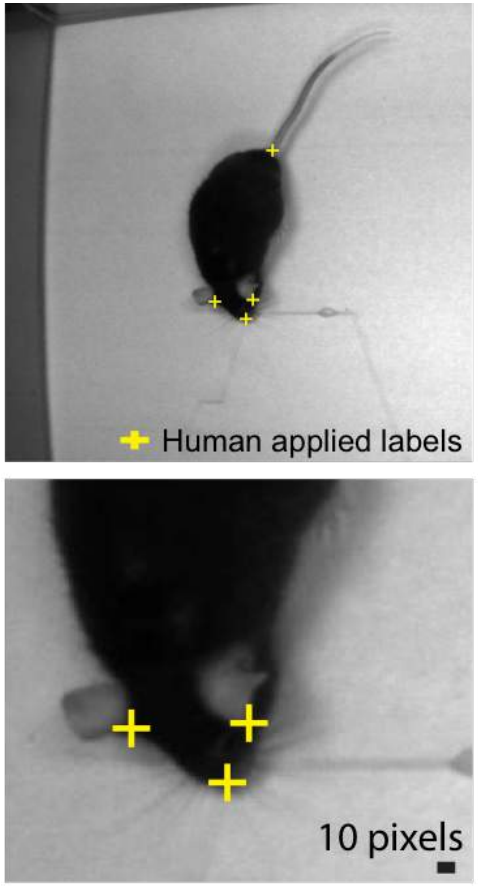

Benchmarking: testing accuracy

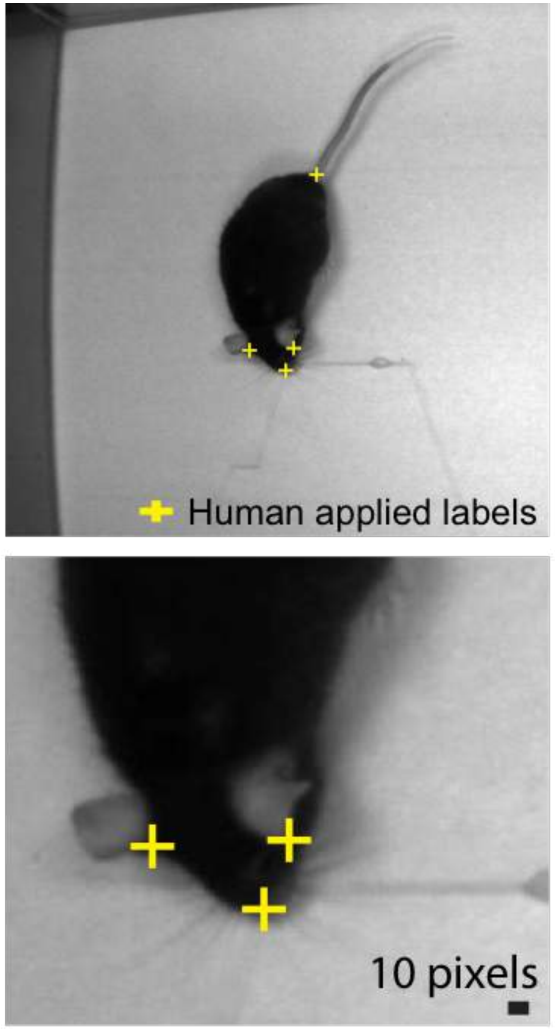

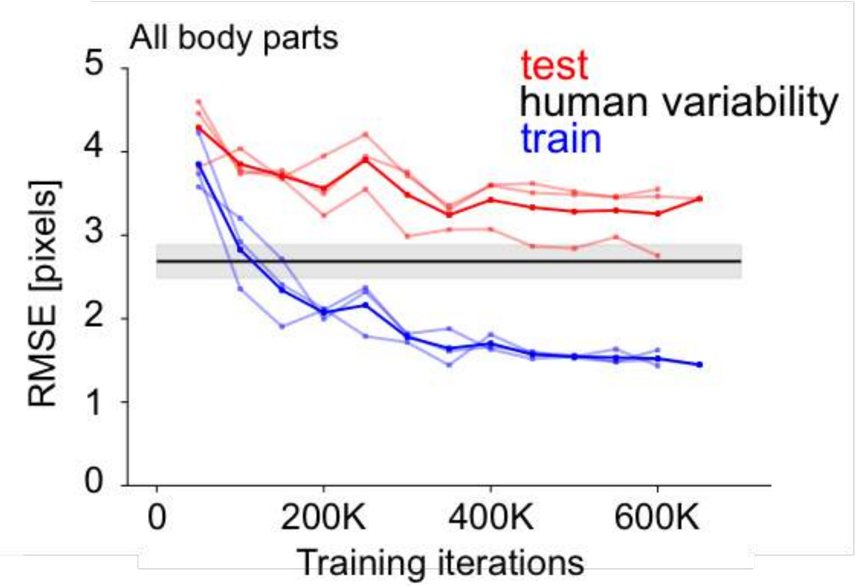

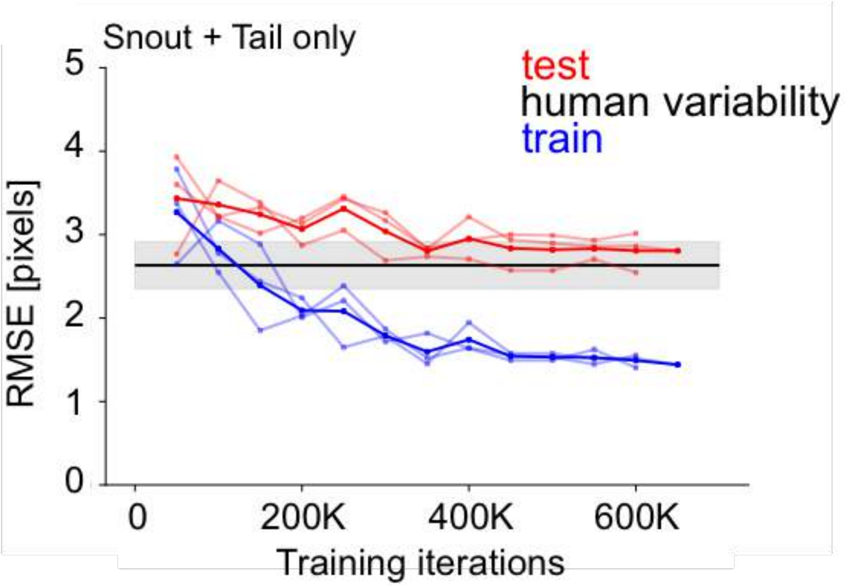

End-to-end training with more labels improves performance

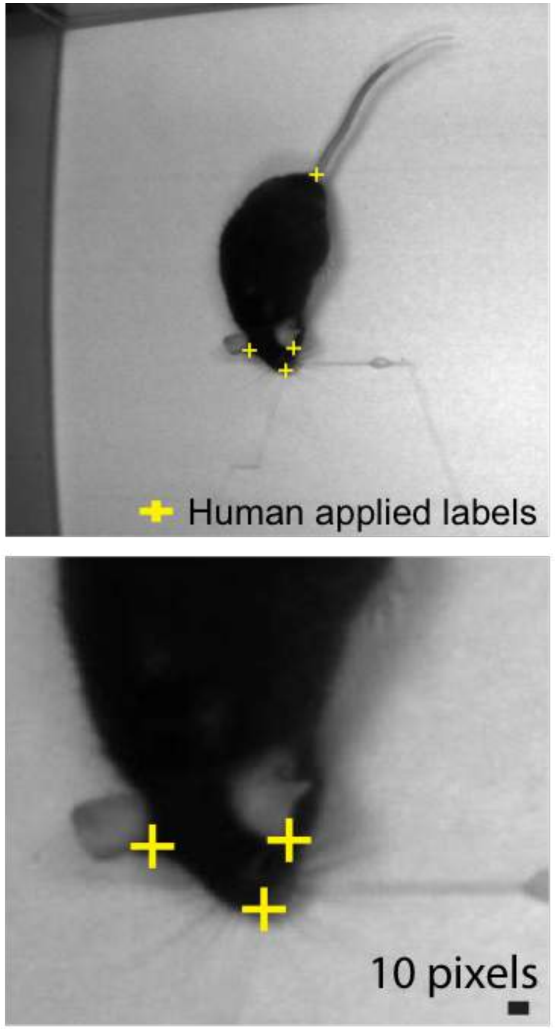

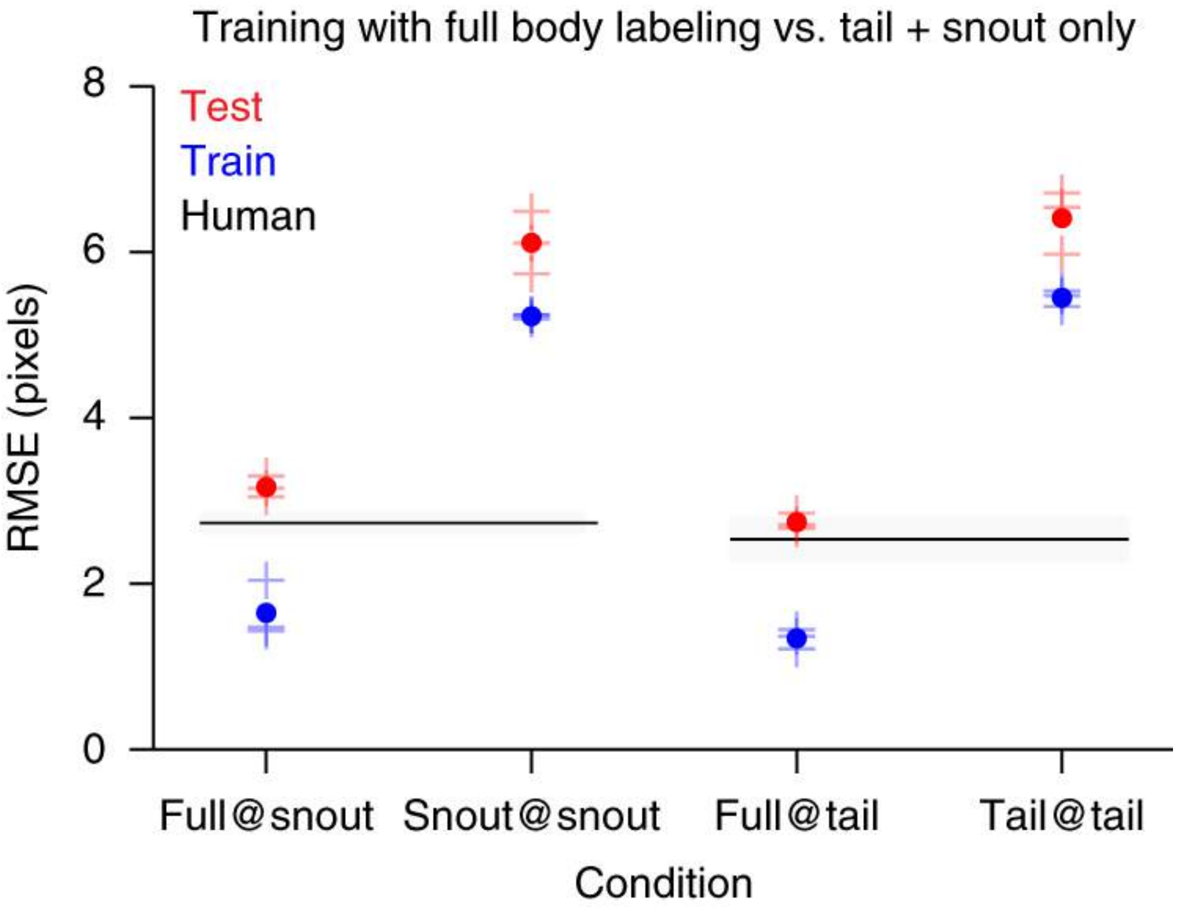

Generalization: multi-mice

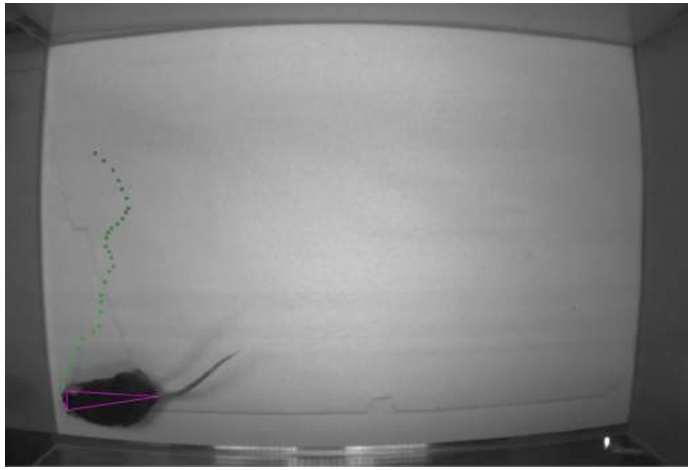

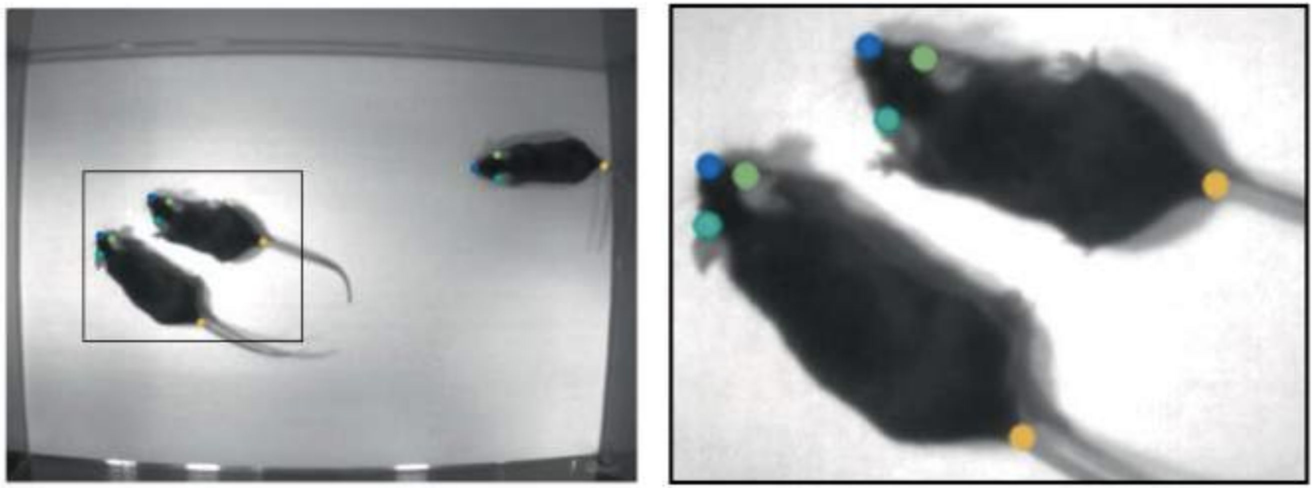

Mathis …. M.W. Mathis* and Bethge*
Nature Neurosci, 2018

Score-maps provide network confidence readout

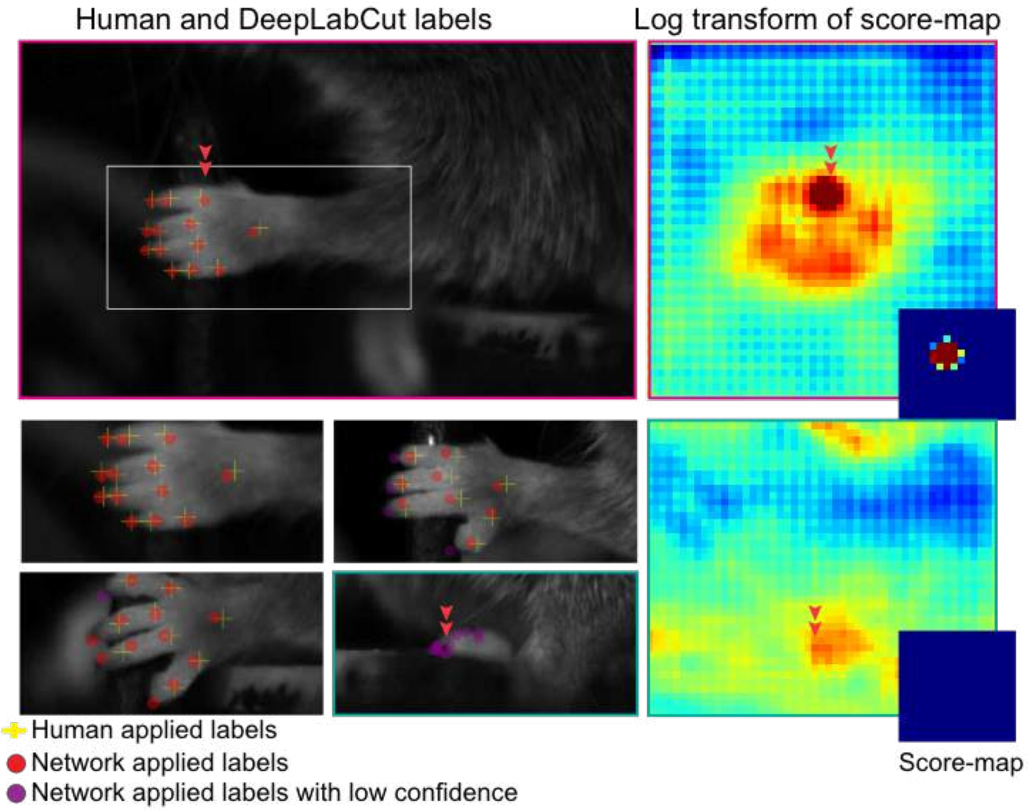

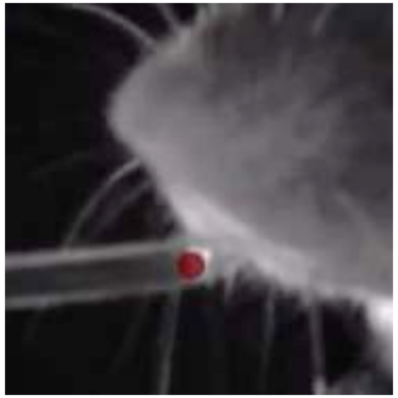

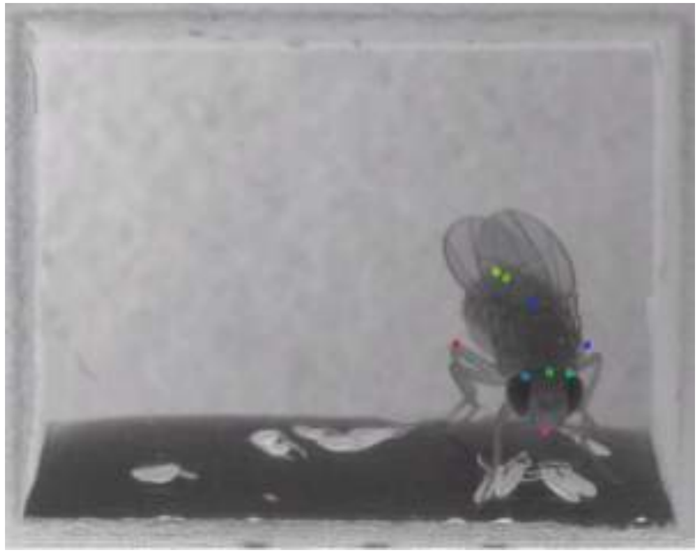

A. Mathis et al, Nature Neuroscience 2018

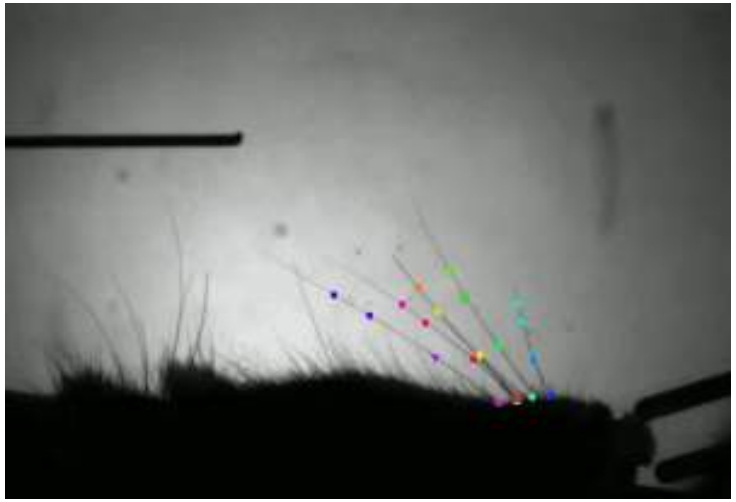

A. Erskine (Hires Lab)

J. Saunders

DeepLabCut Horse Network

11 frames added….

→

Frame-by-frame predictions (no filtering)

Transfer learning enables deep learning in the lab

Pre-trained! (i.e. on ImageNet)

image ⟶ Predictor ⟶ pose

train

Only a few examples (10-200) for most applications

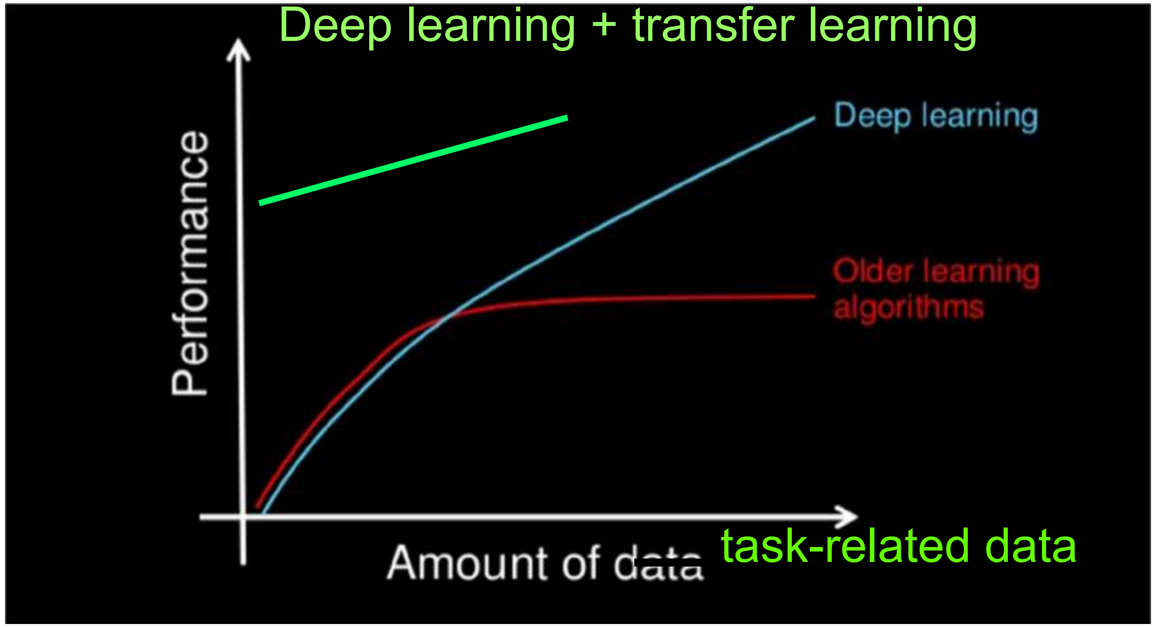

“Software 2.0” – integration of annotation, training and inference

Networks build on ResNets (2018), MobileNetV2 (2019), EfficientNets (2020)

Create a project, extract frames, + GUIs to label your data

Select + Train your deep neural network

Evaluate network performance

(active learning + GUIs if improvement needed)

Run inference on new videos, create labeled videos, + plot your results!

refine?

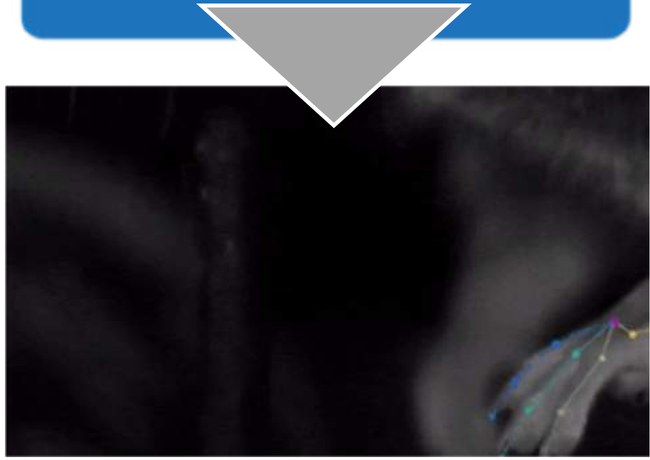

Mathis et al. Nature Neurosci, 2018
Nath* & Mathis* et al. Nature Protocols 2019

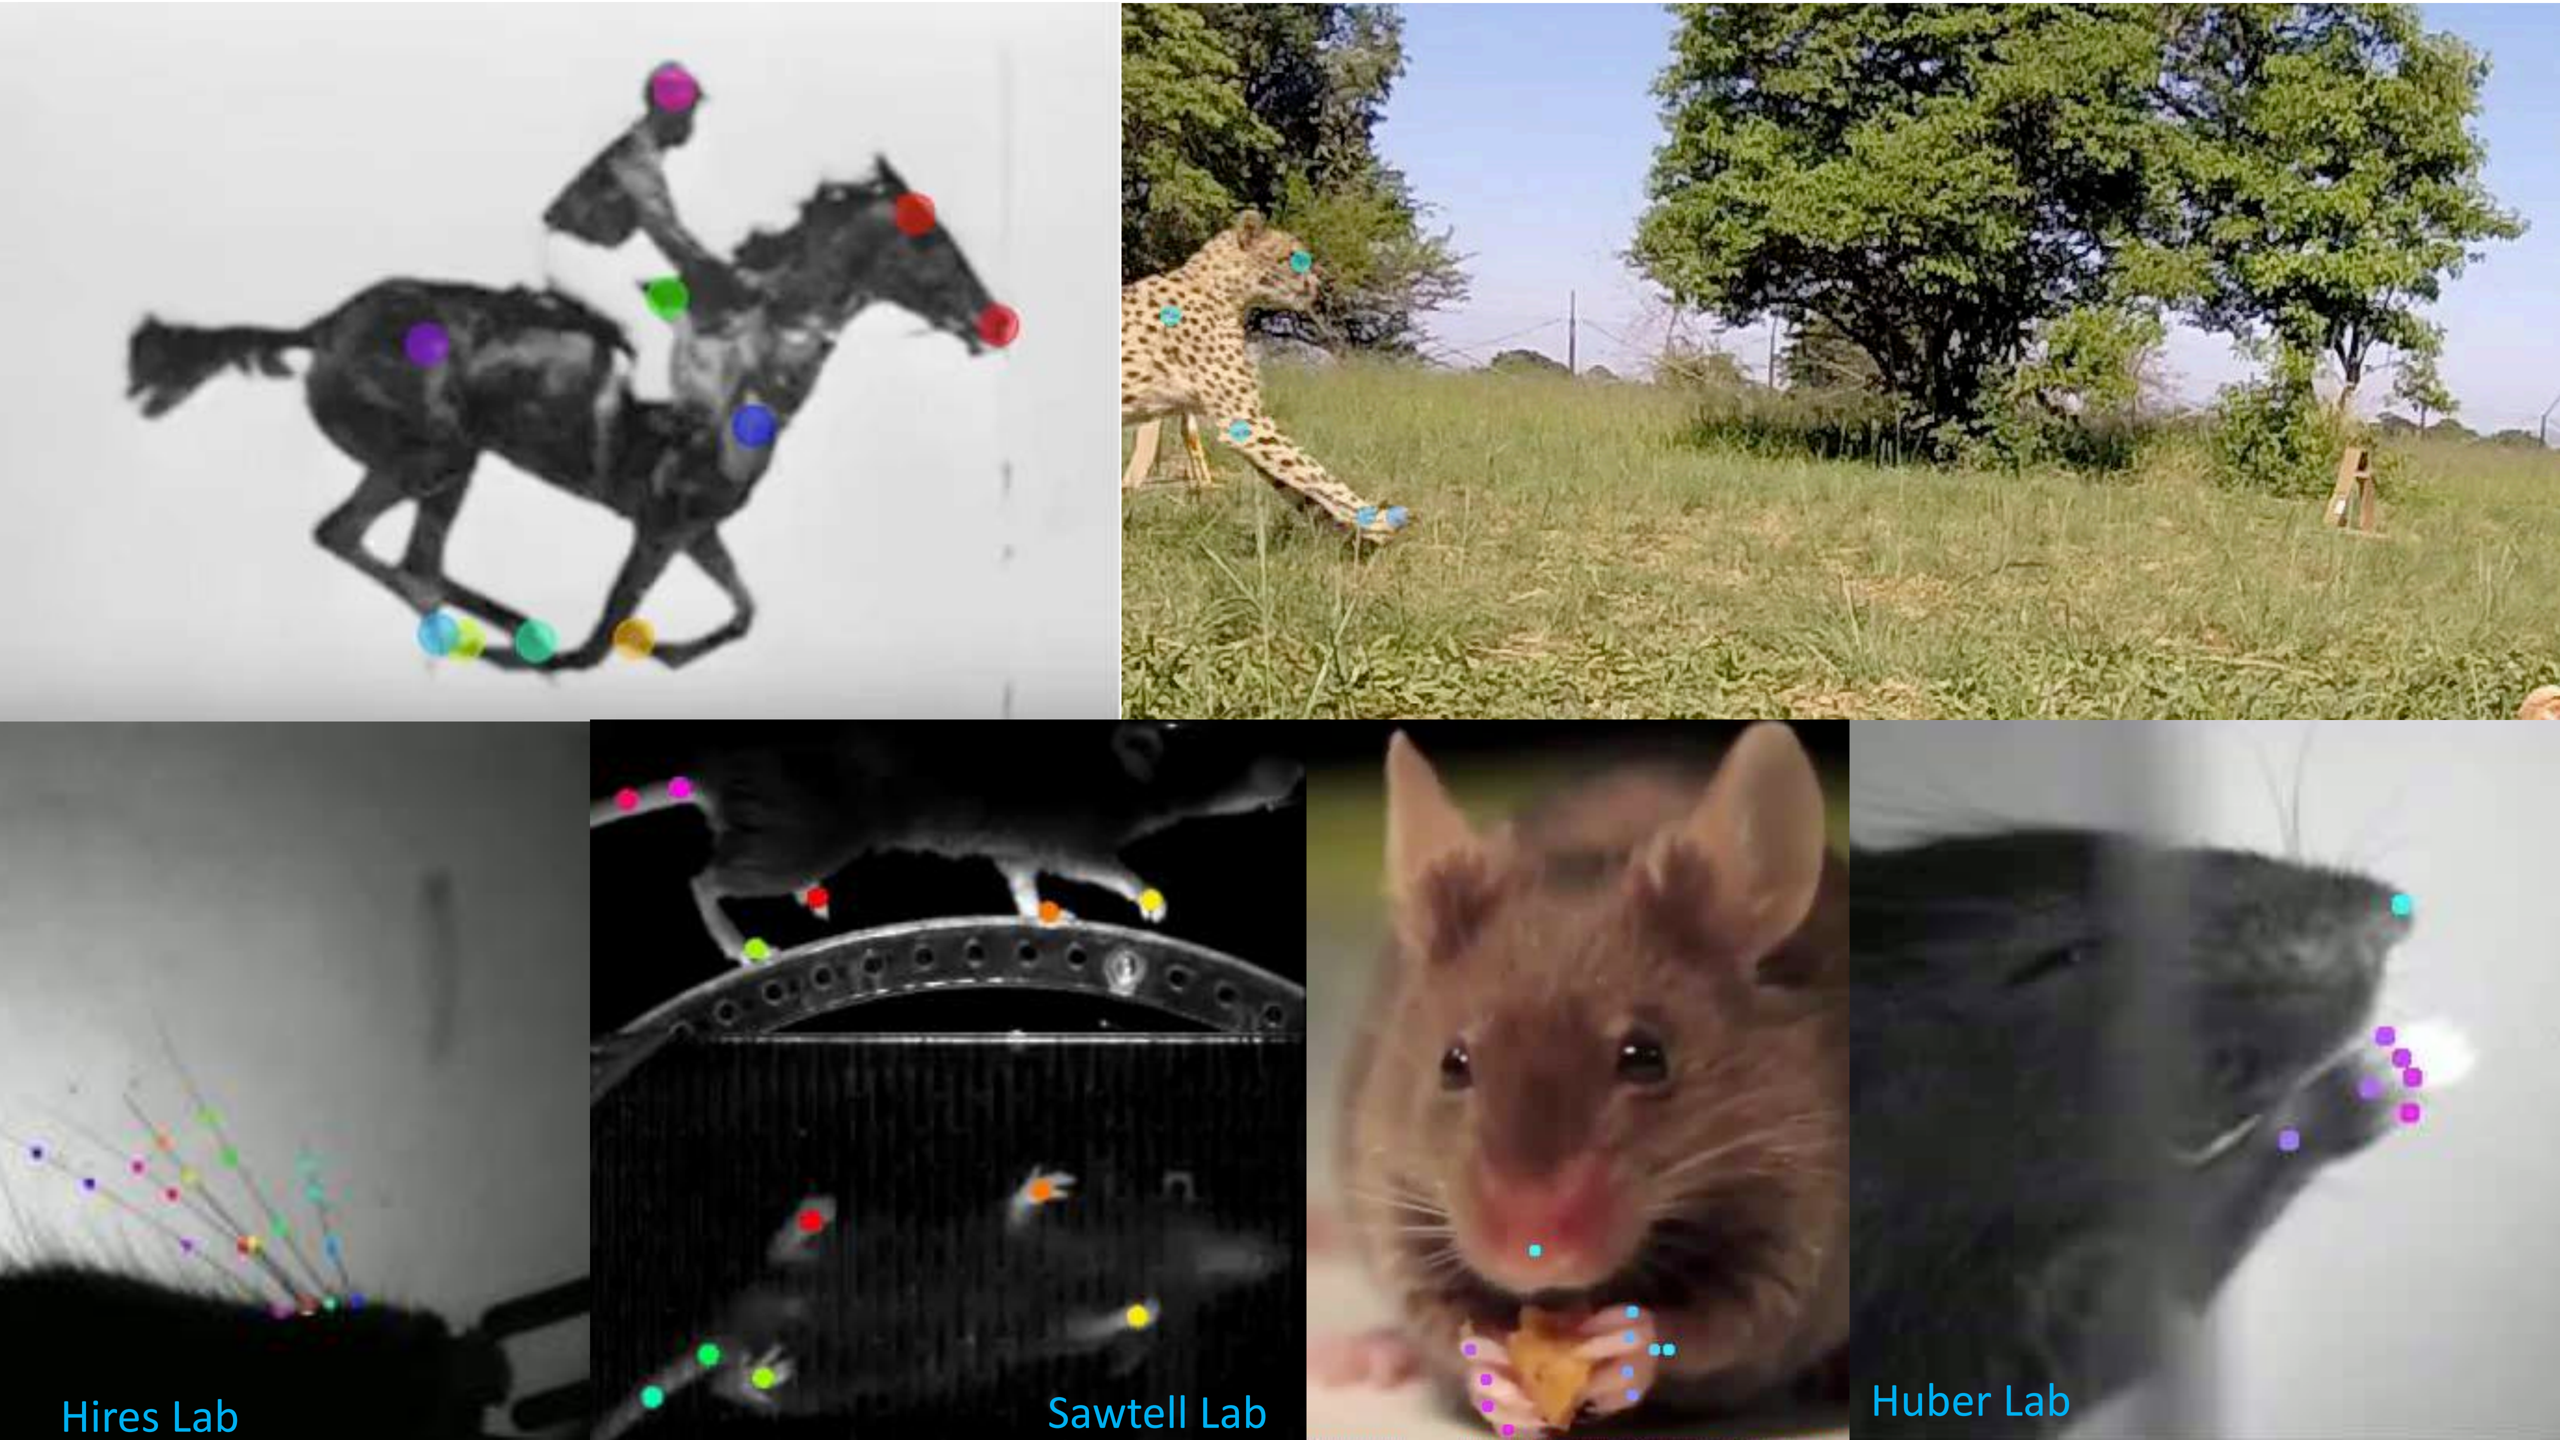
Hires Lab
Sawtell Lab
Huber Lab

DeepLabCut: training **one** network for 3D tracking of a hunting cheetah

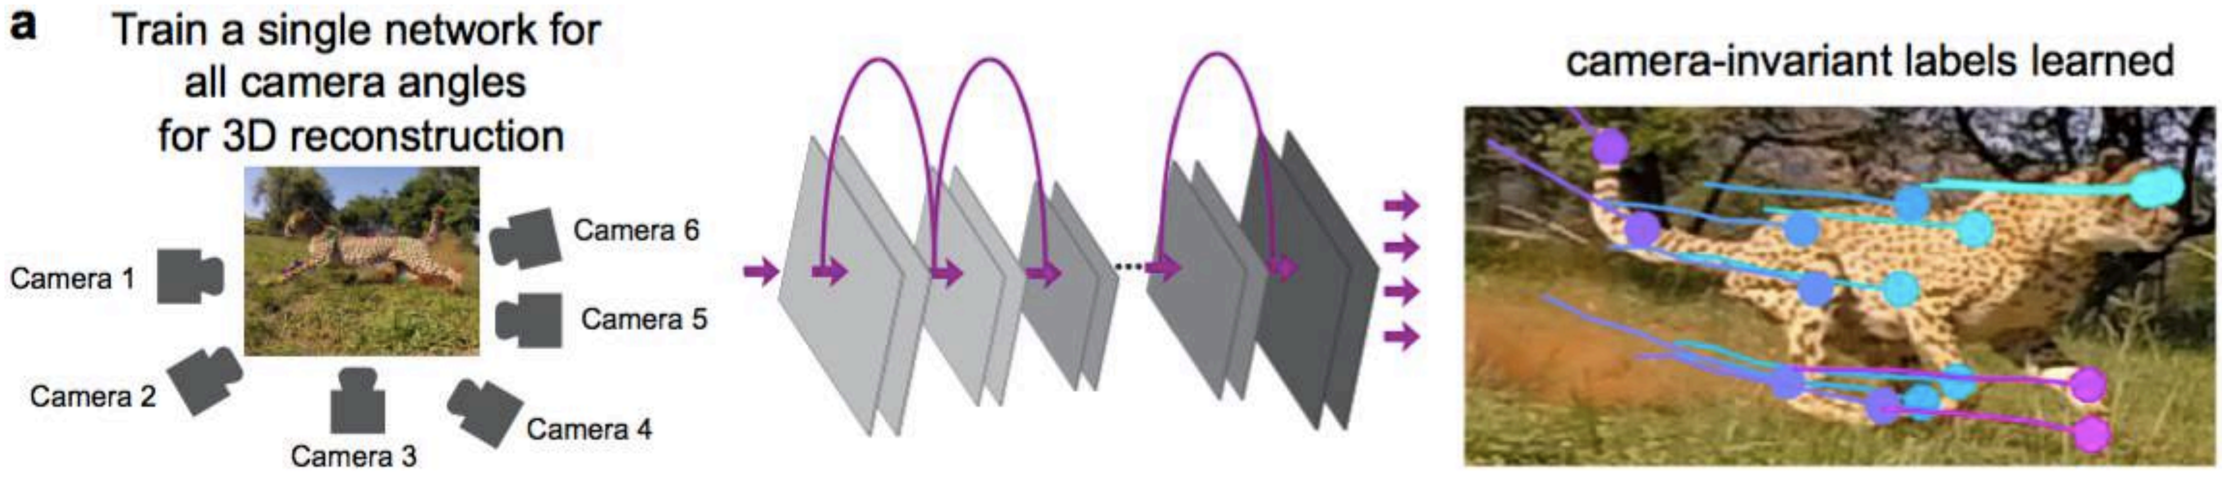

Nath*, Mathis* et al. Nature Protocols (in press)

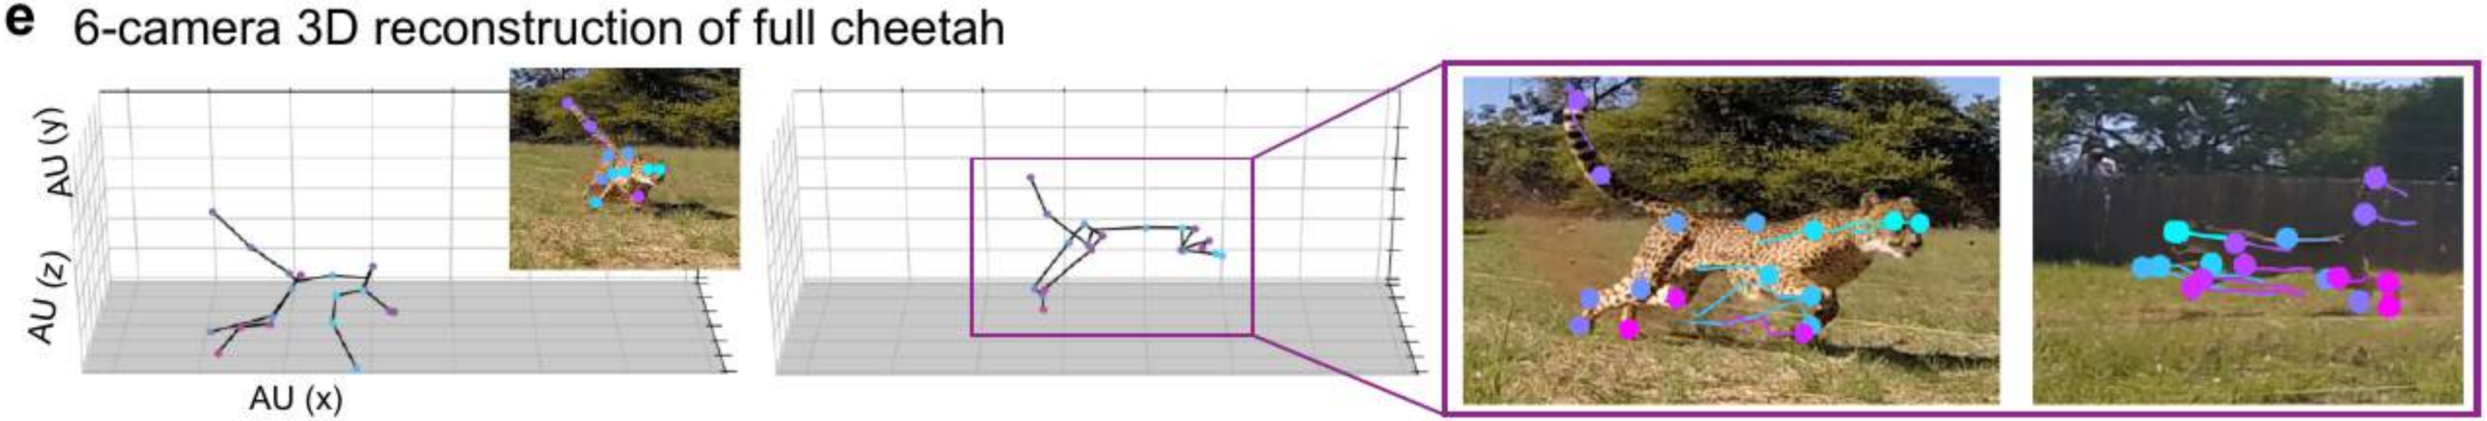

Nath*, Mathis* et al. Nature Protocols (in press)

3D pose estimation

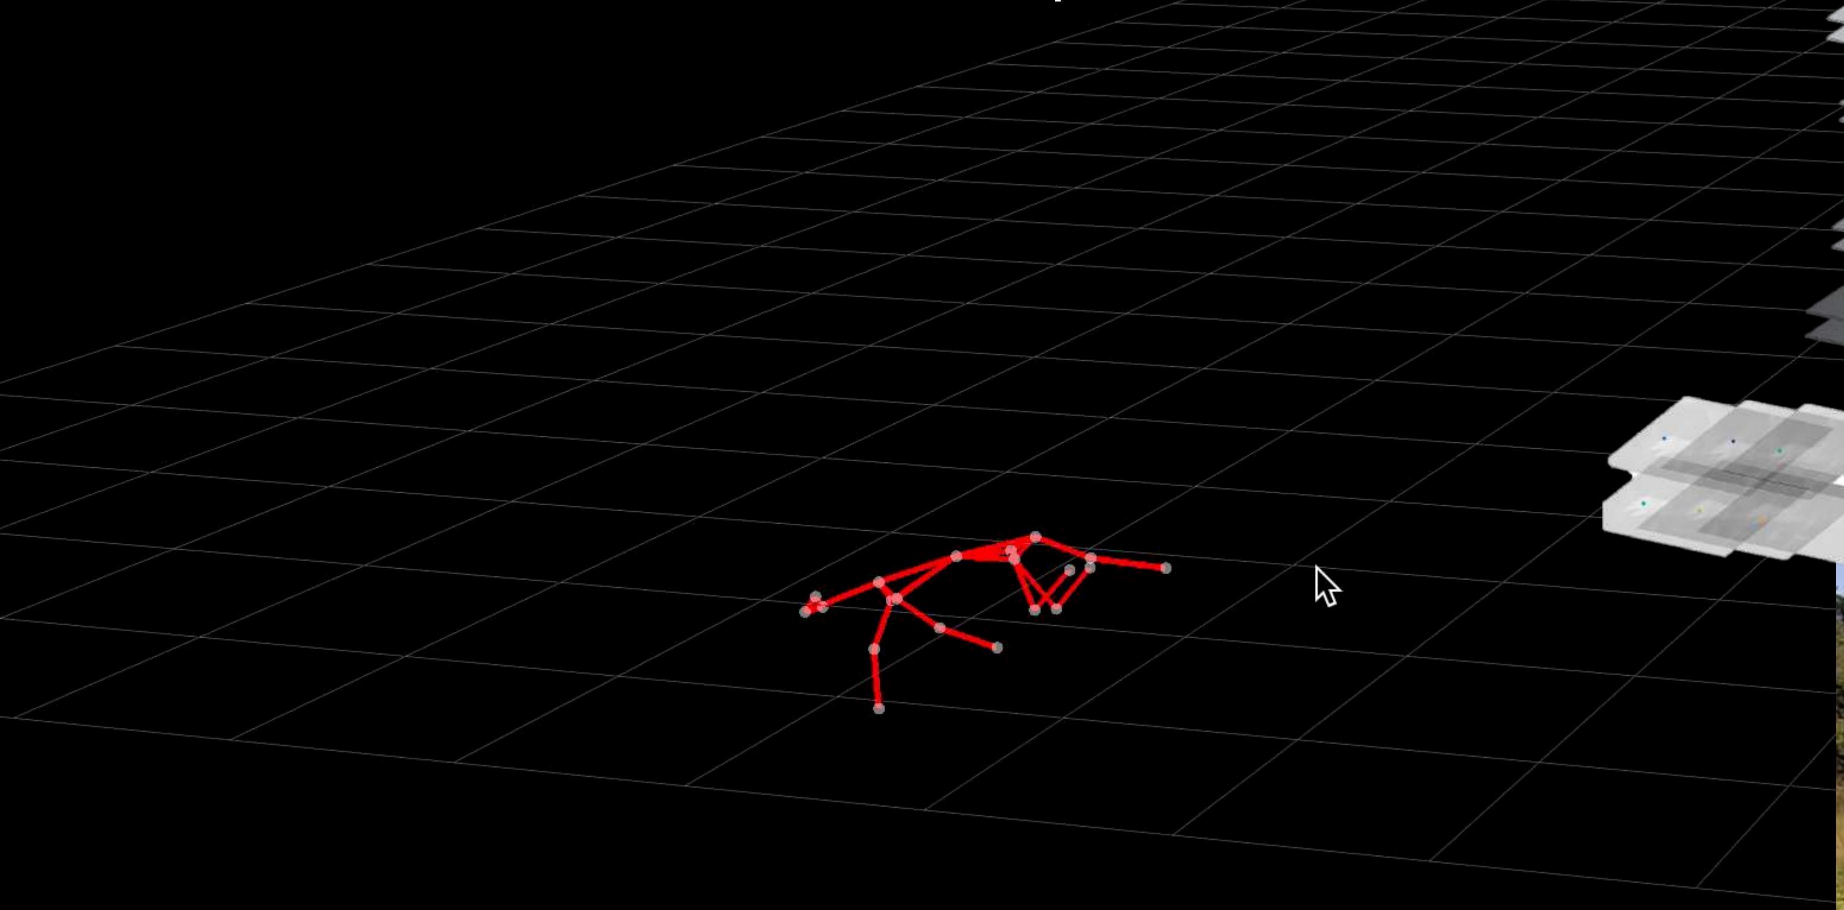

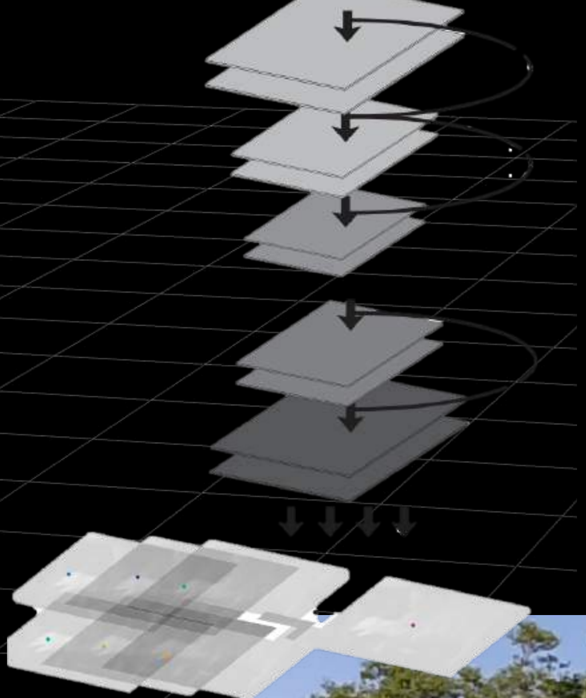

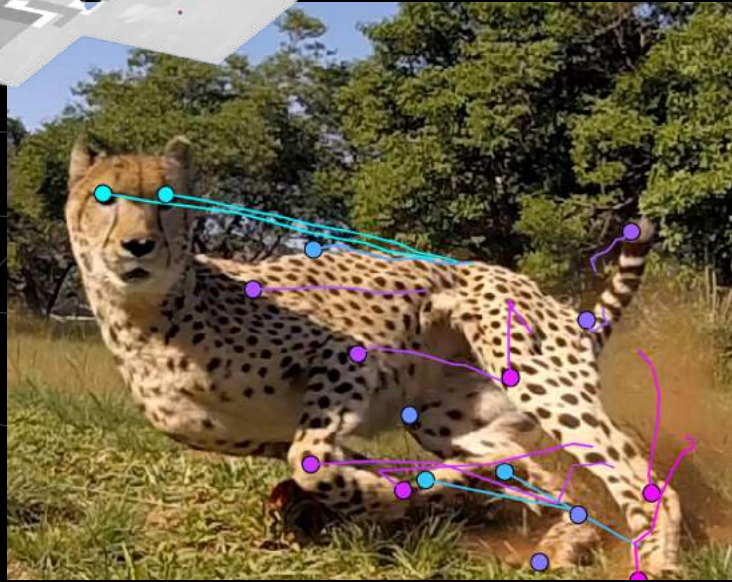

Courtesy Amir Patel (U of Cape Town)
Nath et al. Nat. Protocols, 2019

How does markerless pose estimation work?

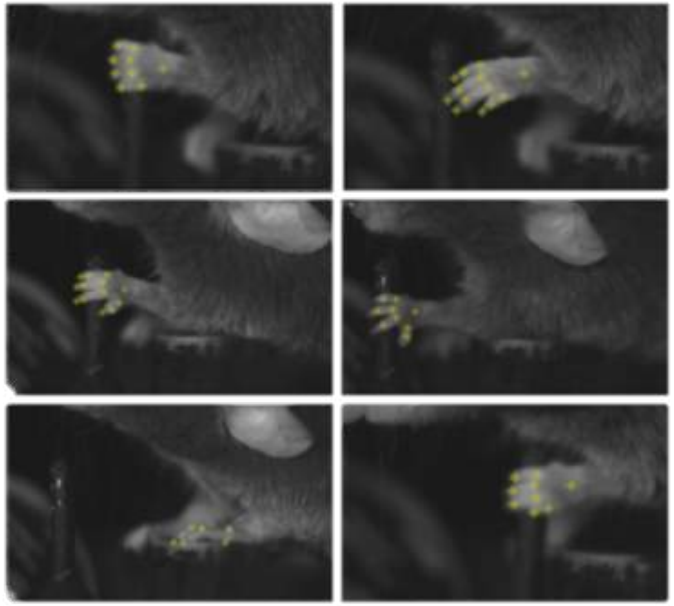

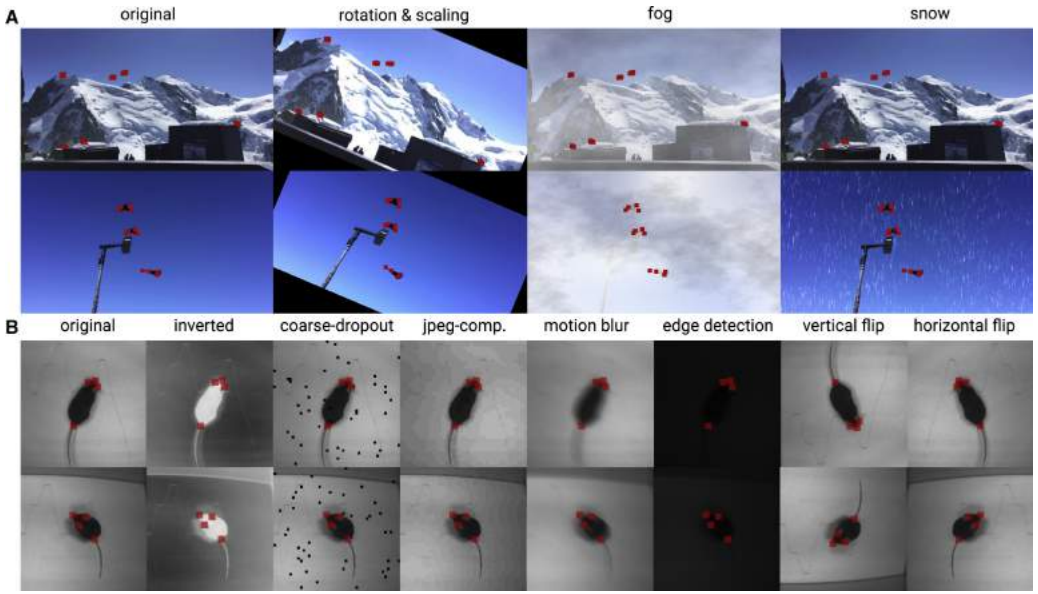

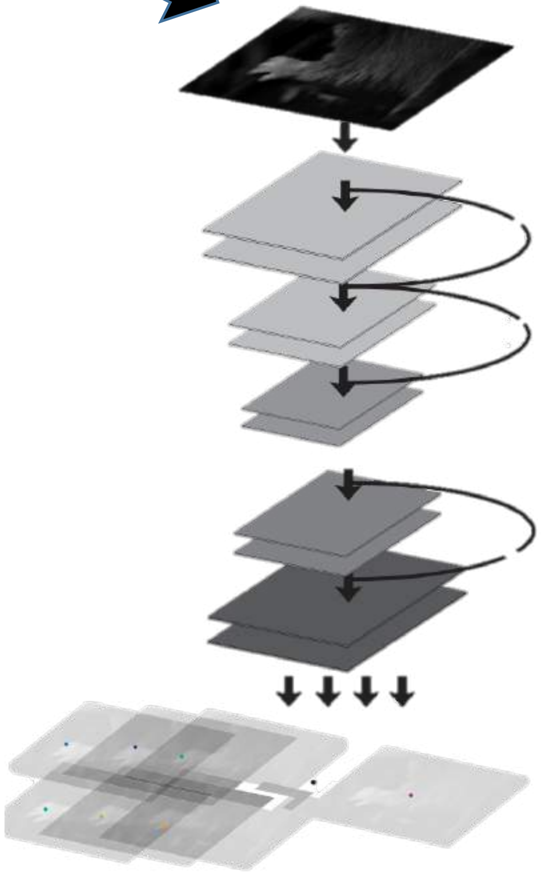

Key Features:
- Data augmentation
- Model architecture
- Optimization

Mathis, Schneider, Lauer & Mathis, 2020 Neuron

Data Augmentation: how to get the most out of your data!

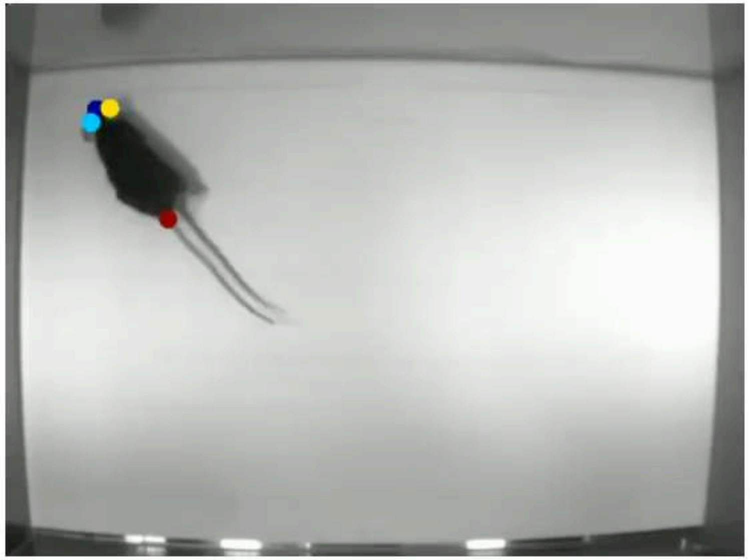

tensorpack

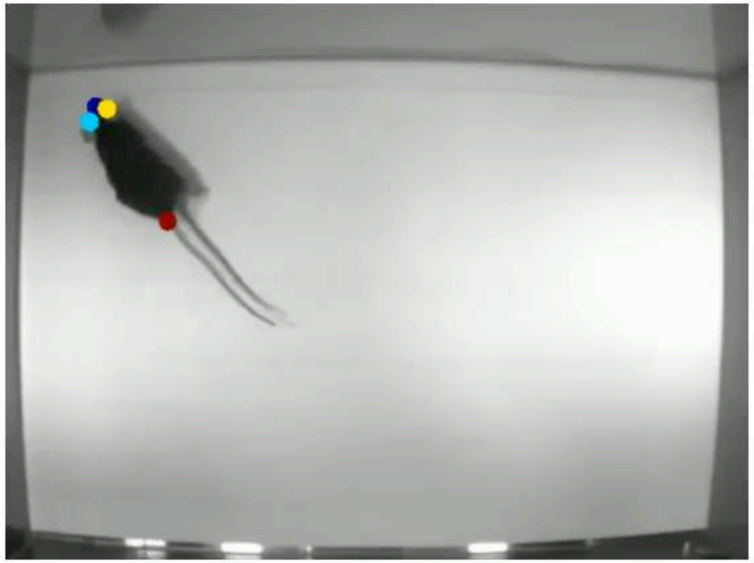

imgaug

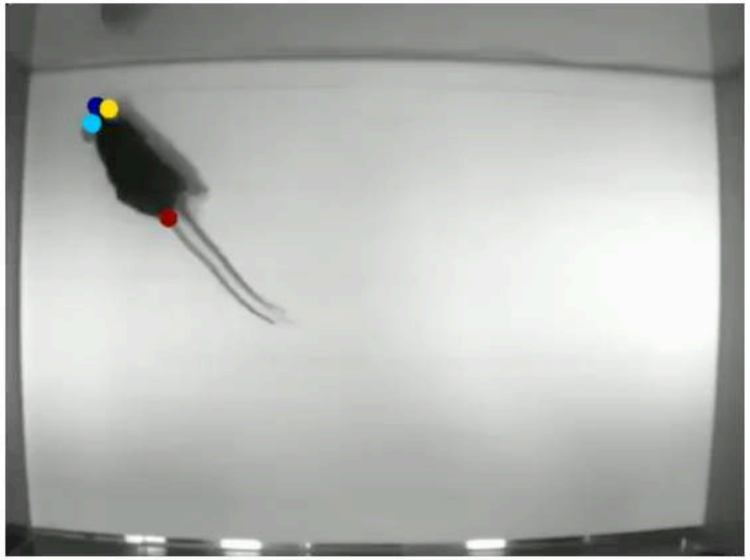

scale-crop

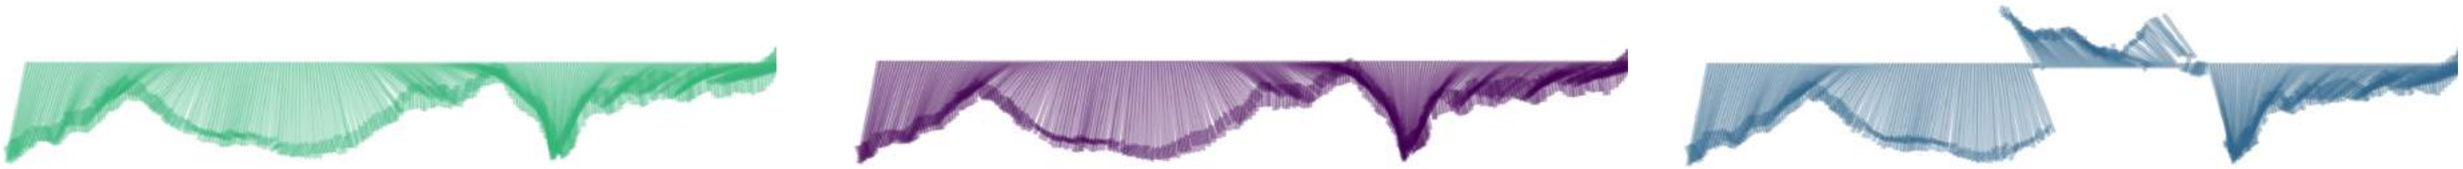

Mathis, Schneider, Lauer & Mathis
2020 Neuron

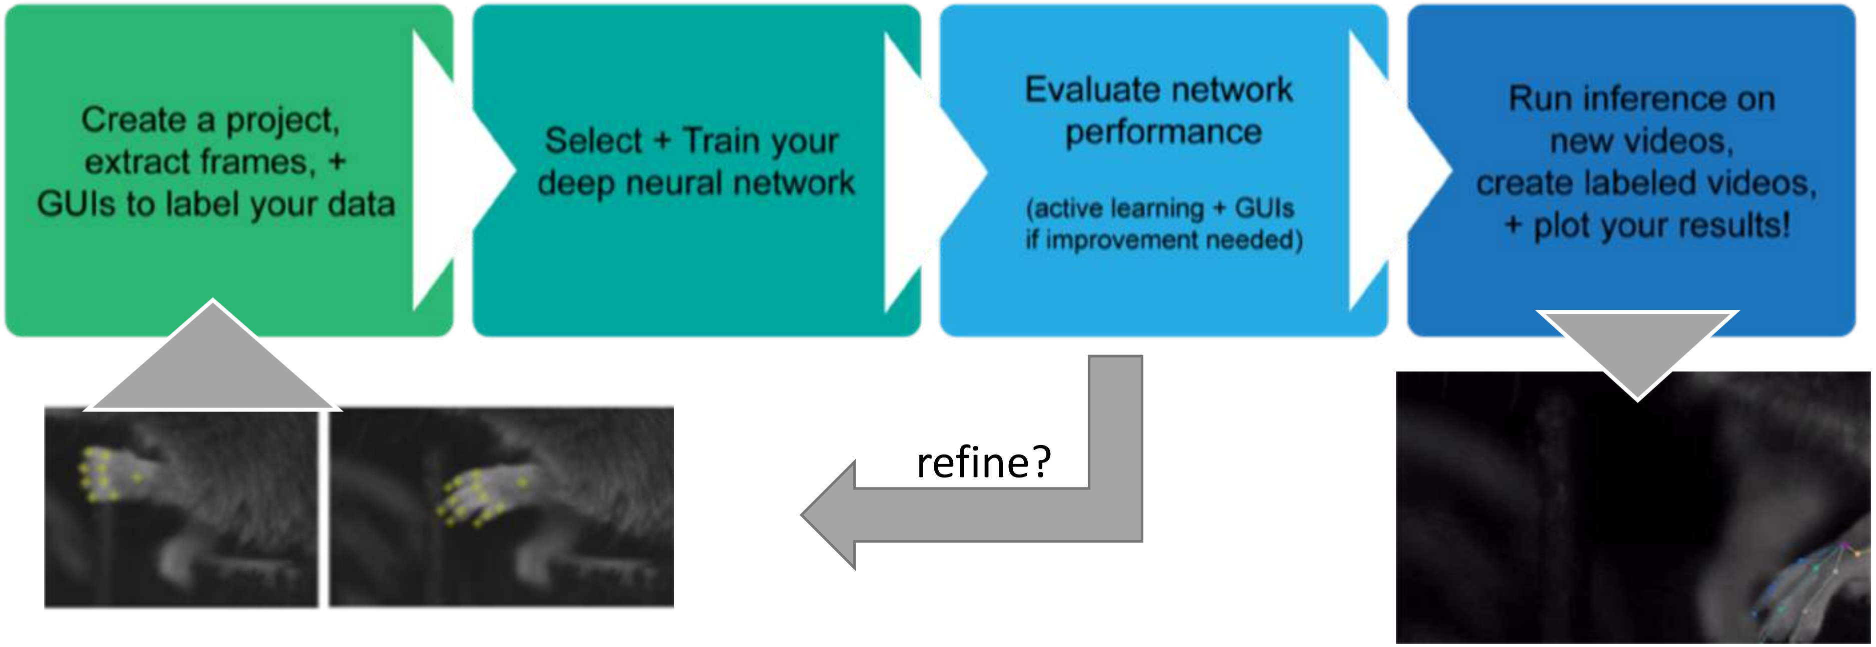

Mathis et al 2020 Neuron

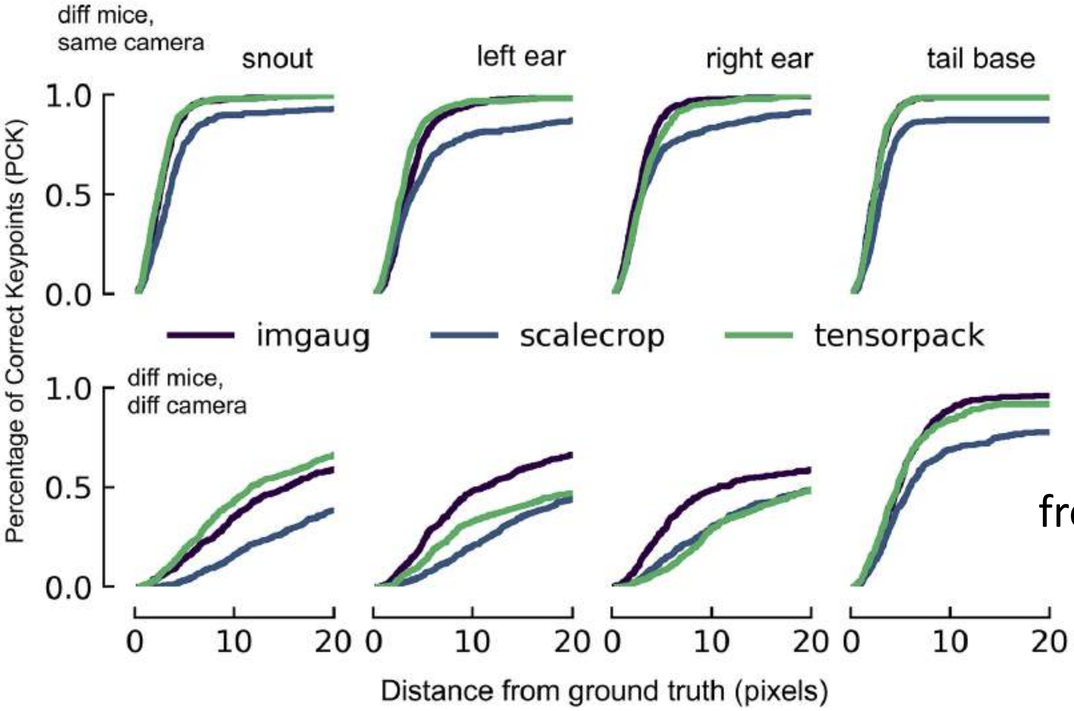

refine:

28 frames added from another camera

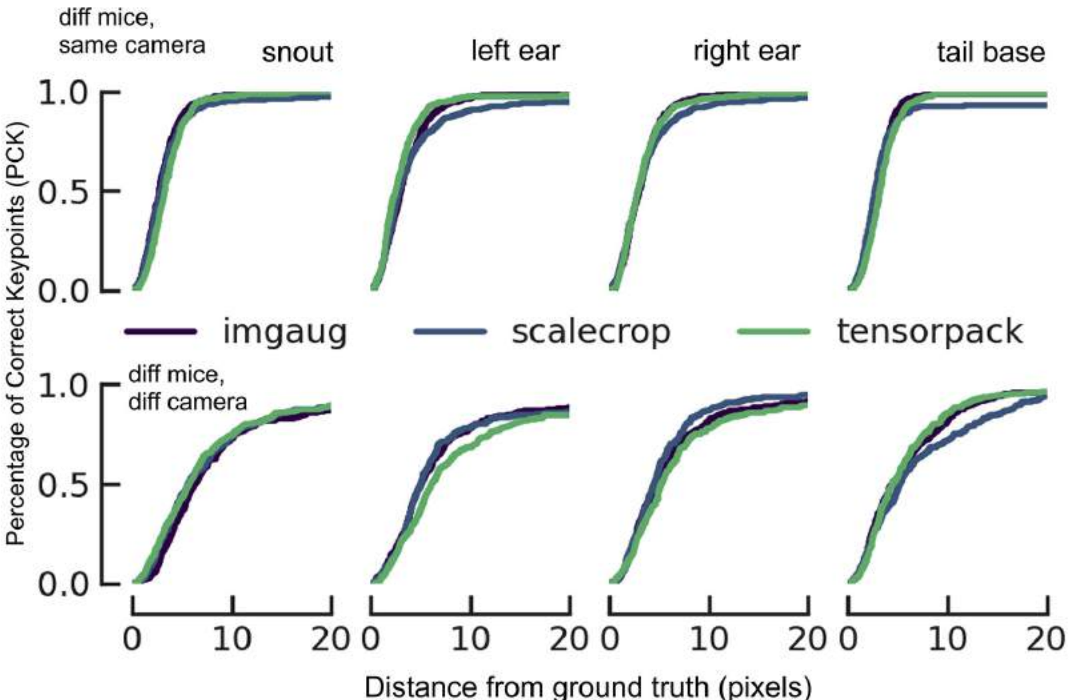

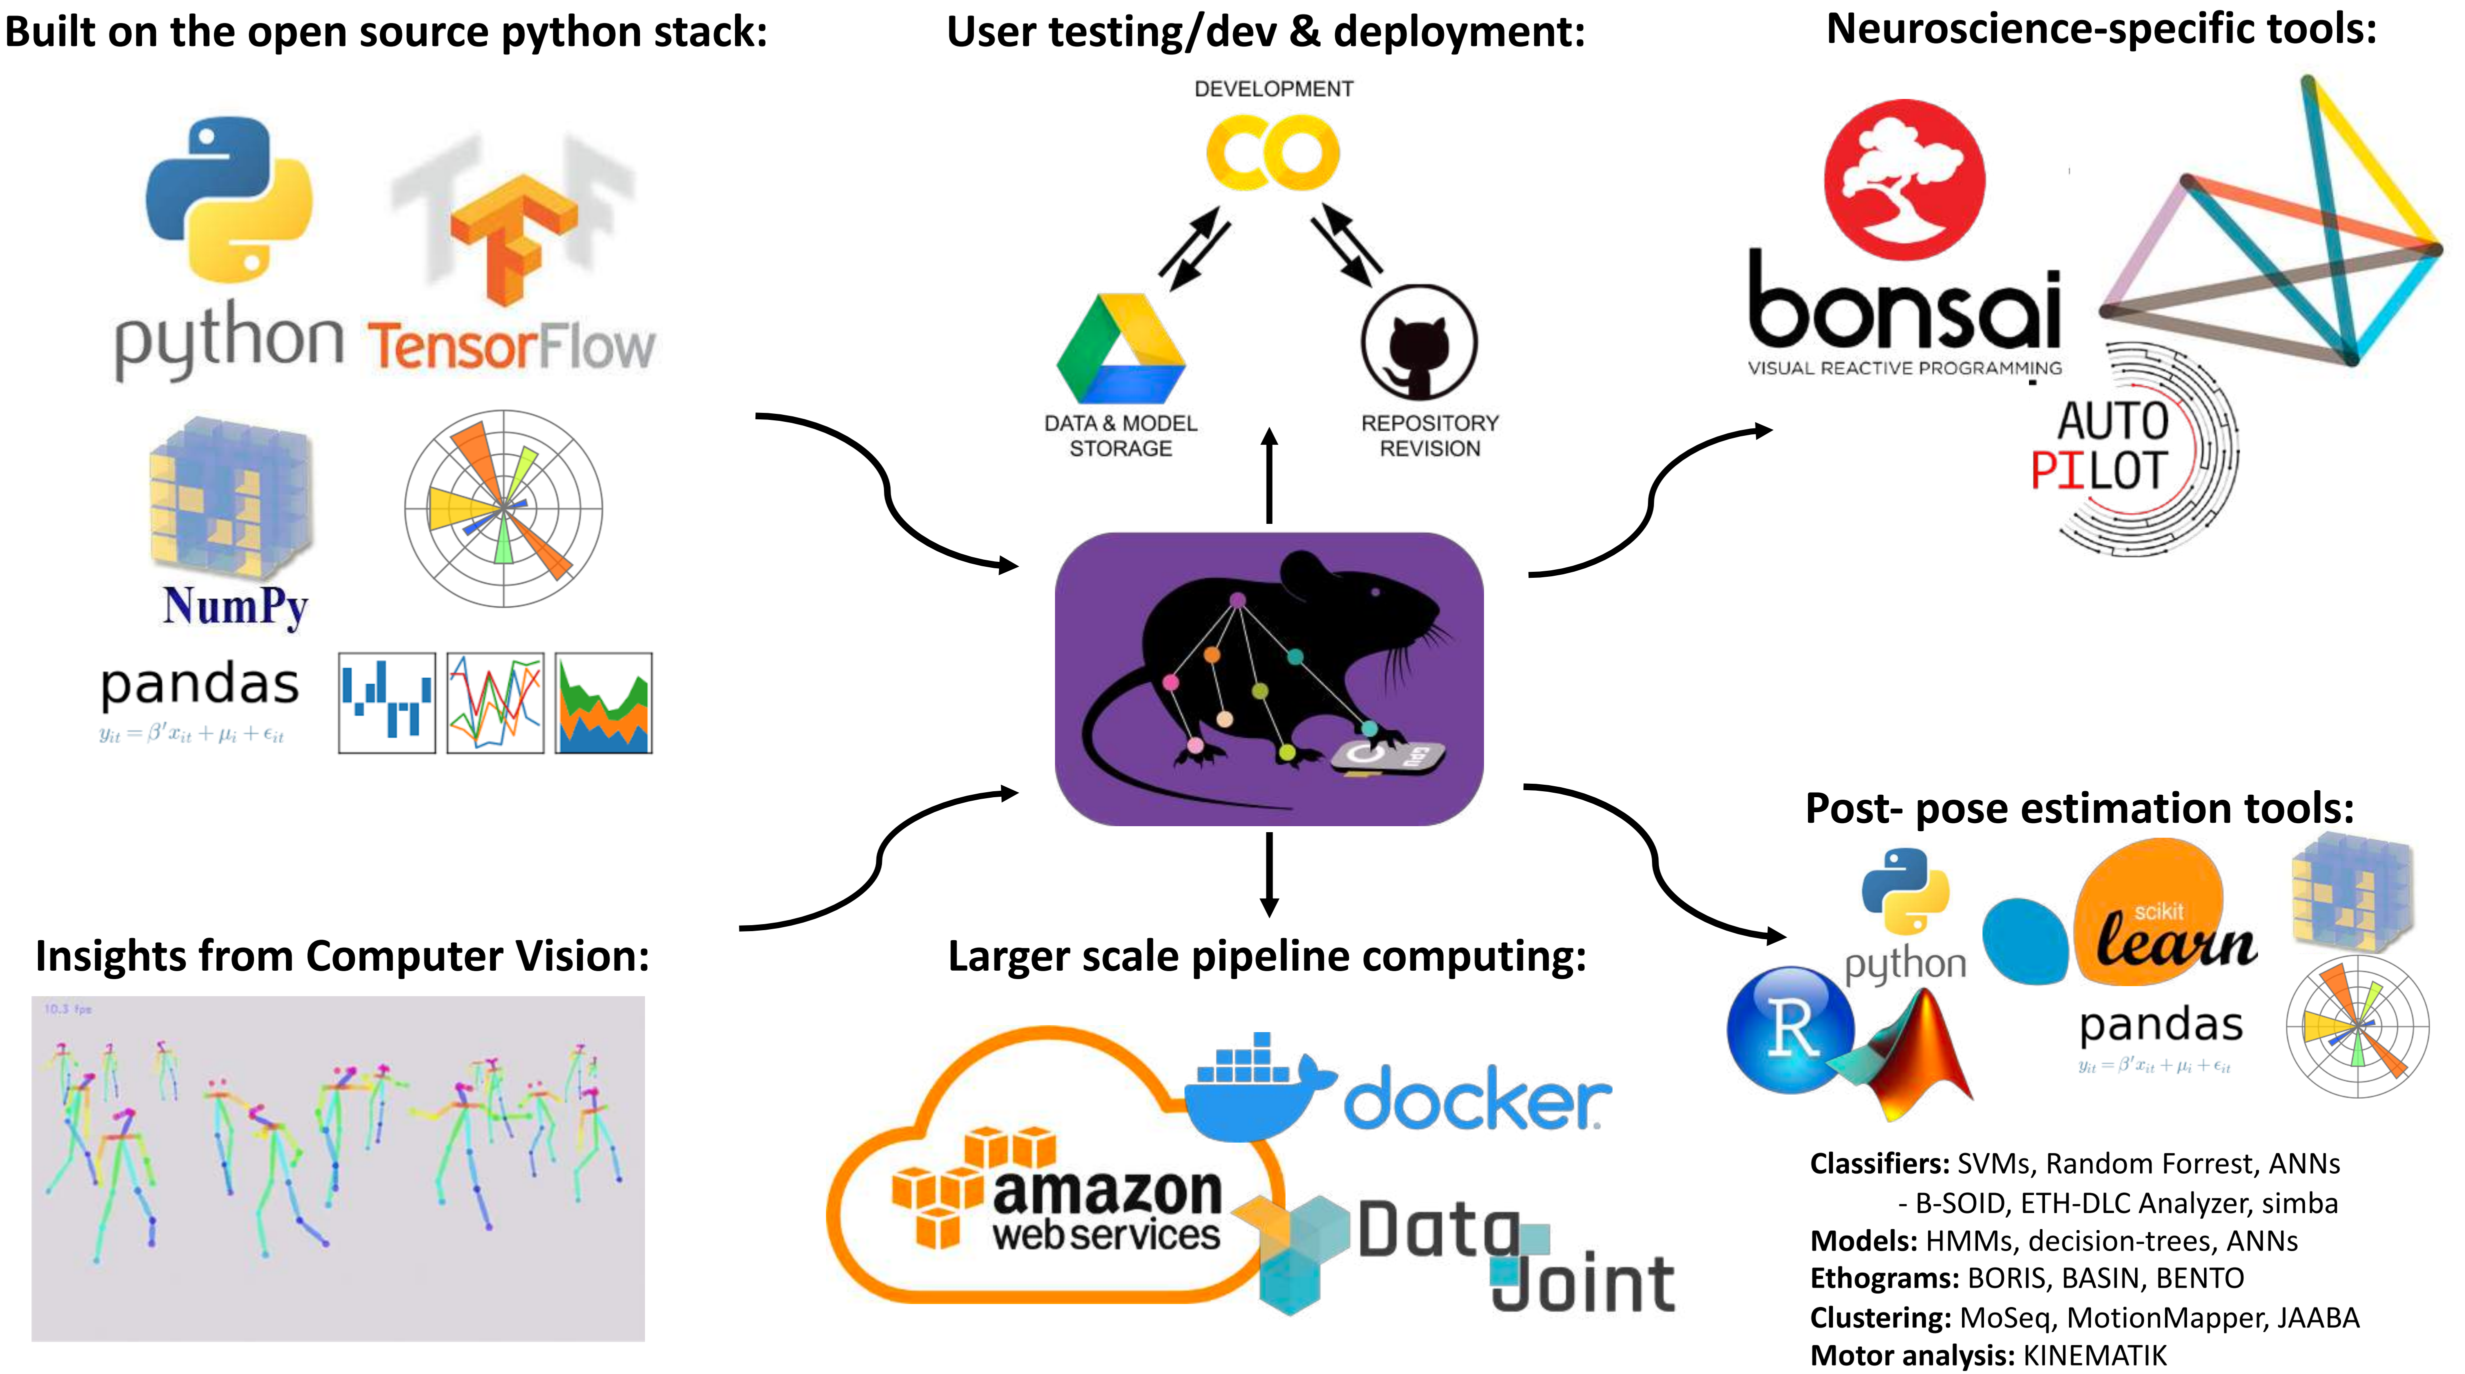
Built on the open source python stack:
python
TensorFlow
NumPy
pandas
$y_{it} = \beta' x_{it} + \mu_i + \epsilon_{it}$
User testing/dev & deployment:
DEVELOPMENT
CO
DATA & MODEL STORAGE
REPOSITORY REVISION
Neuroscience-specific tools:
bonsai
VISUAL REACTIVE PROGRAMMING
AUTO PILOT
Insights from Computer Vision:
Larger scale pipeline computing:
amazon web services
docker
DataJoint
Post- pose estimation tools:
python
scikit learn
R
pandas
$y_{it} = \beta' x_{it} + \mu_i + \epsilon_{it}$
Classifiers: SVMs, Random Forrest, ANNs
- B-SOID, ETH-DLC Analyzer, simba
Models: HMMs, decision-trees, ANNs
Ethograms: BORIS, BASIN, BENTO
Clustering: MoSeq, MotionMapper, JAABA
Motor analysis: KINEMATIK

How AI can influence new behavioral paradigms (and model systems…)

- Robust tracking can allow for more natural behaviors, and better analysis of classical paradigms …

- Opens new research avenues in behavioral phenotyping, new tools for relating neural circuits to behavior …

Natural behaviors

Diverse species, natural environments, many animals

Looming stimulus

M. Meister/Cal Tech

Multi-directional reaching
(and neural reconstruction with DL)

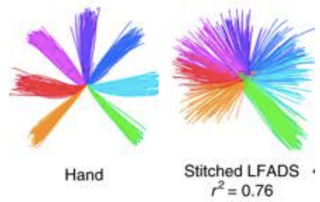

LFADS – Pandarinath et al 2018

No trial structure

Classical conditioning

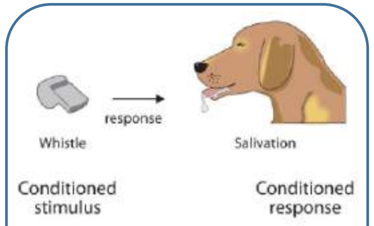

http://open.lib.umn.edu/intropsyc/

Robust & “simple” trial-based behaviors

Complex & more variable (trial-based) behaviors

BREAK ….

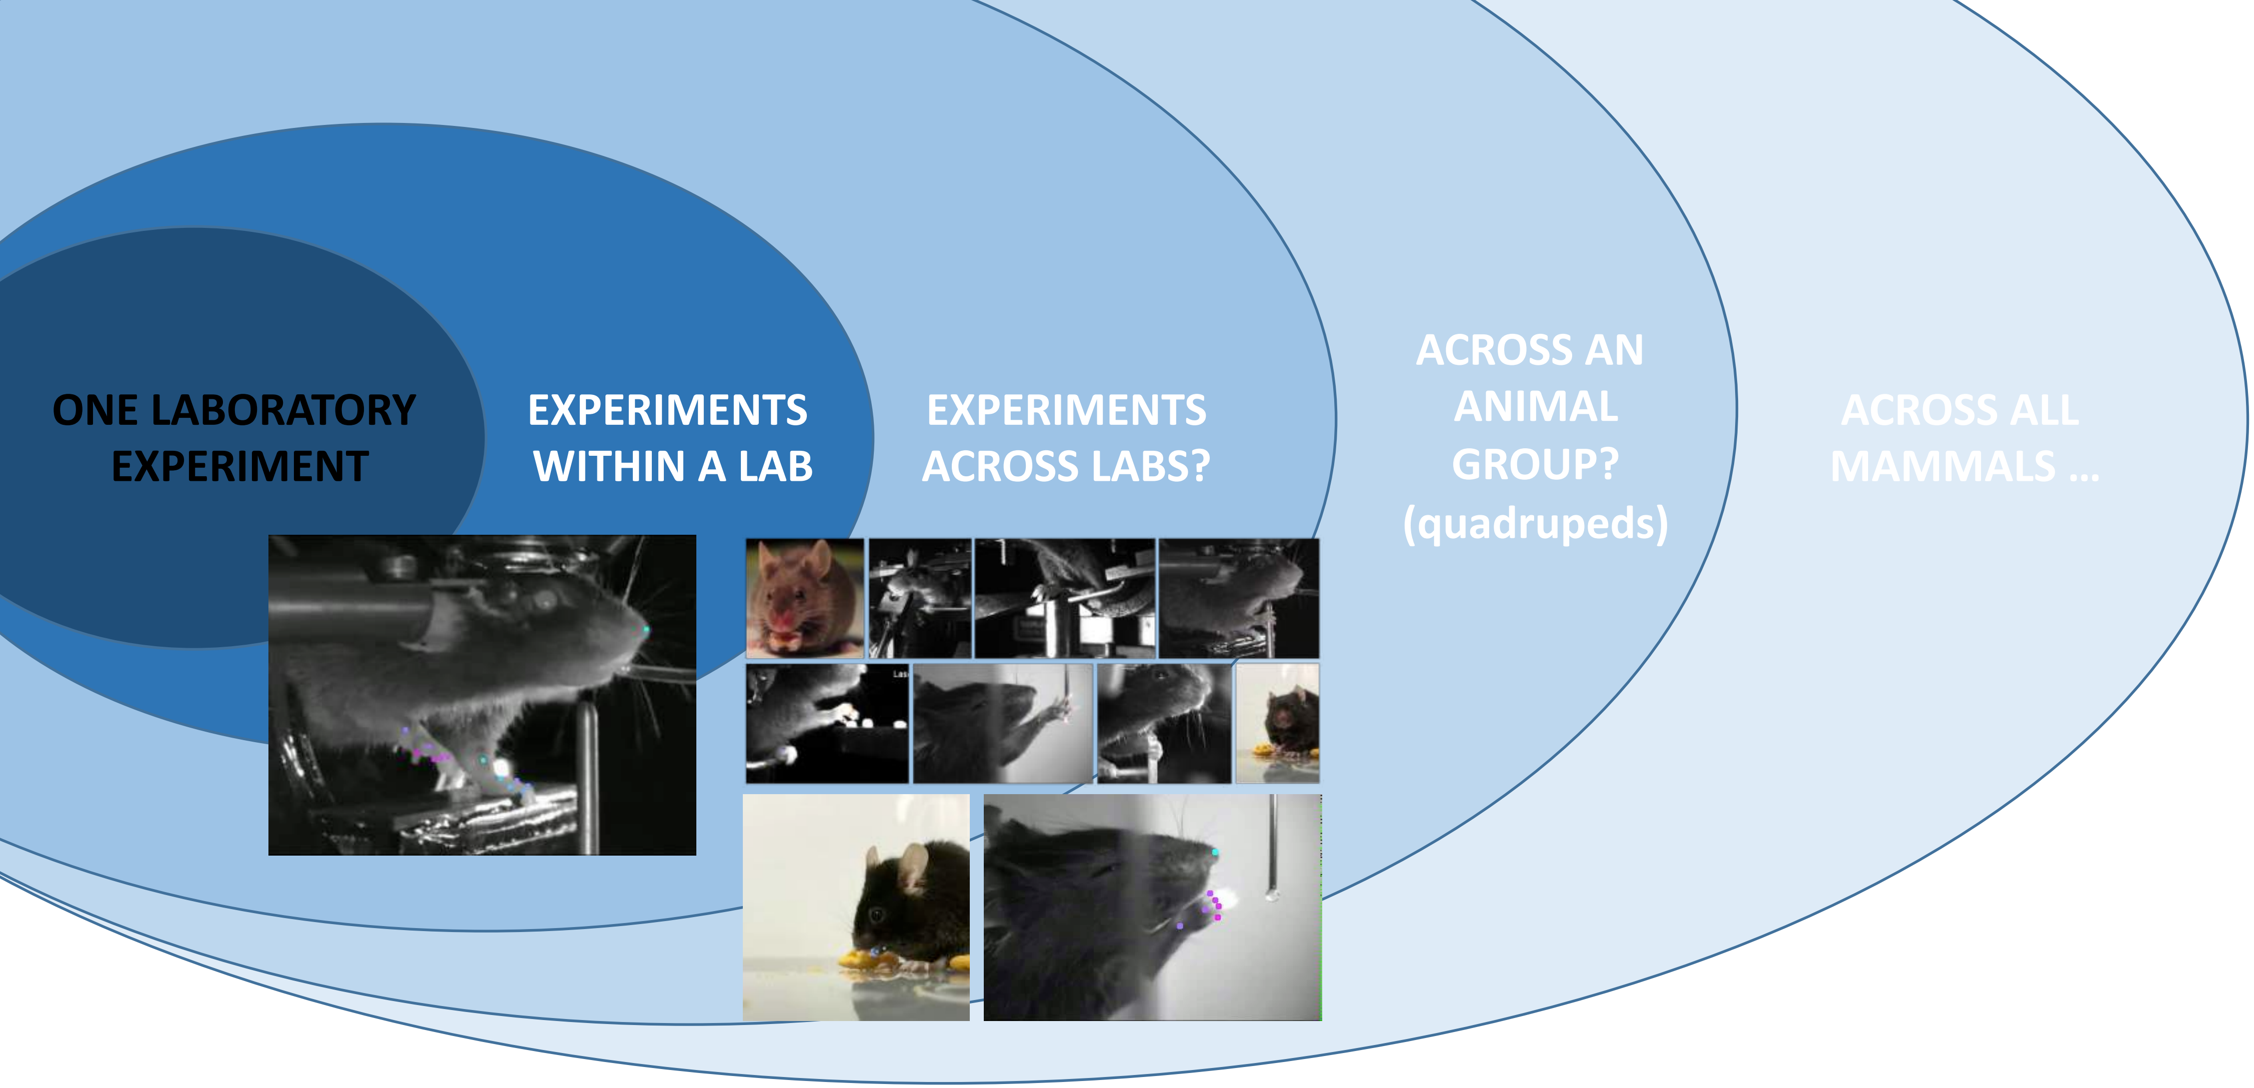

The larger the scale, the greater the diversity ….

How do we create generalizable pose estimation networks?

An animal pose benchmark (horses) for robustness

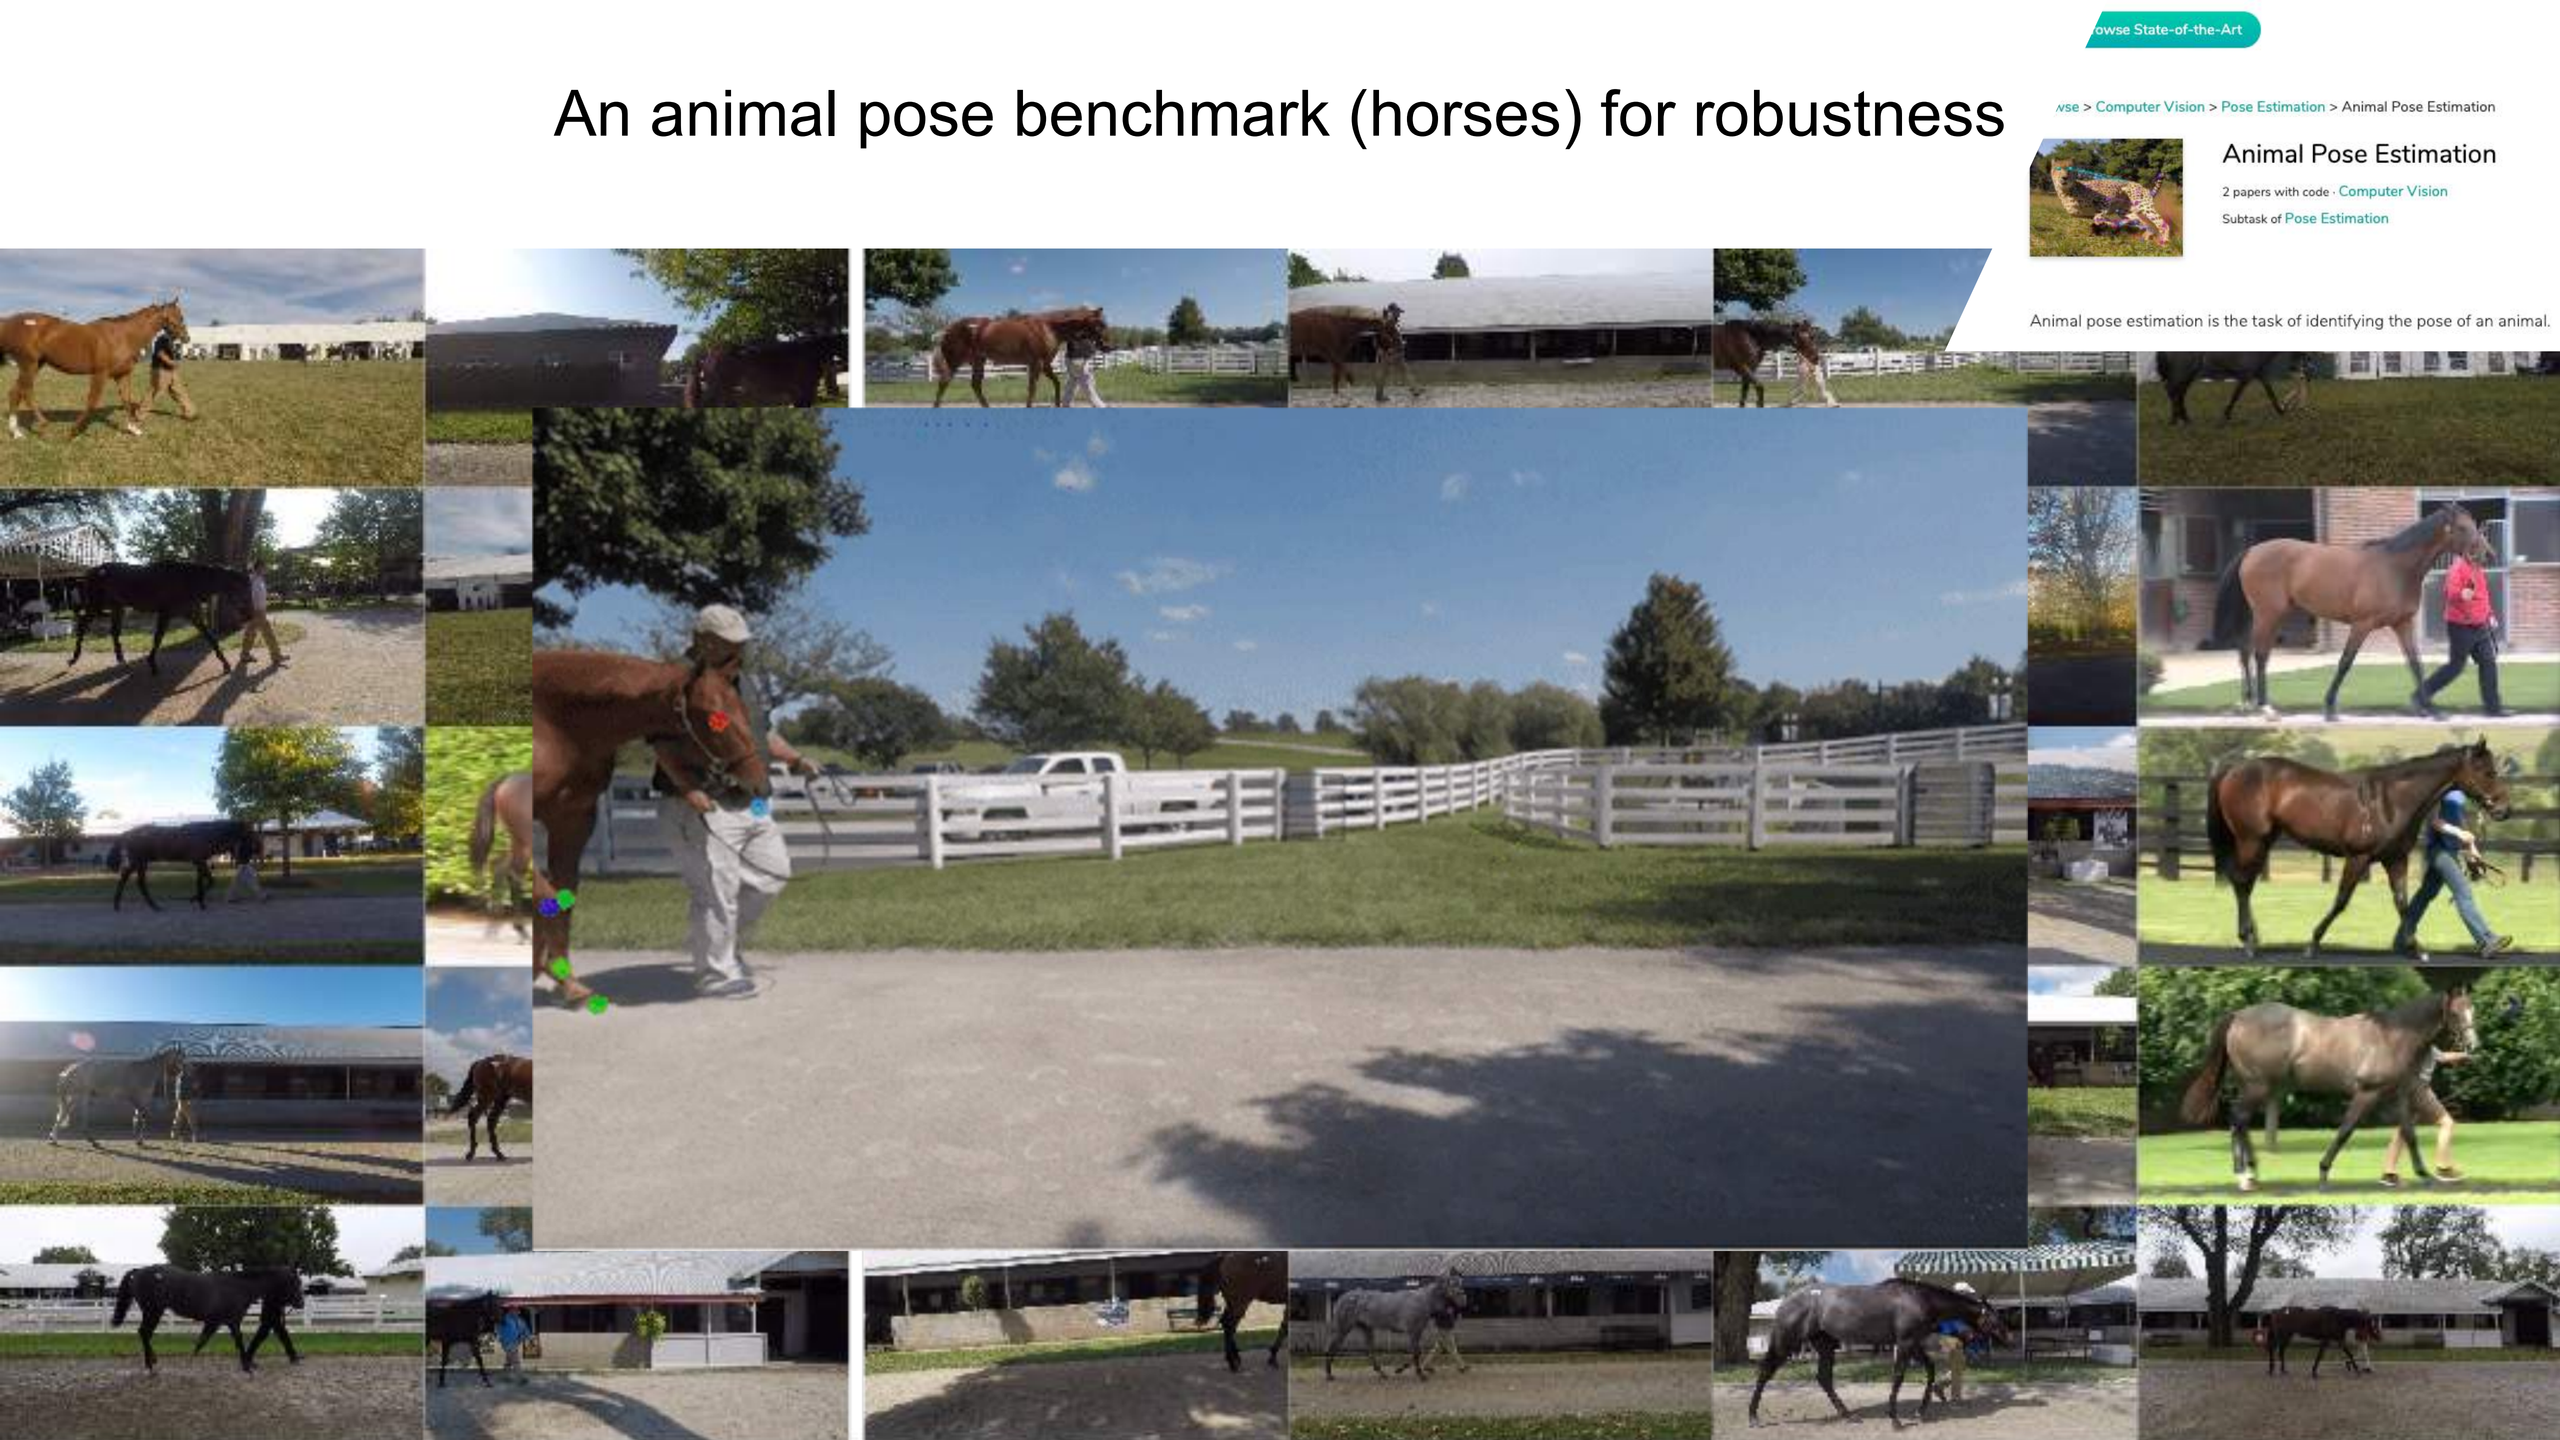

ImageNet performance correlates with pose estimation robustness and generalization on out-of-domain data

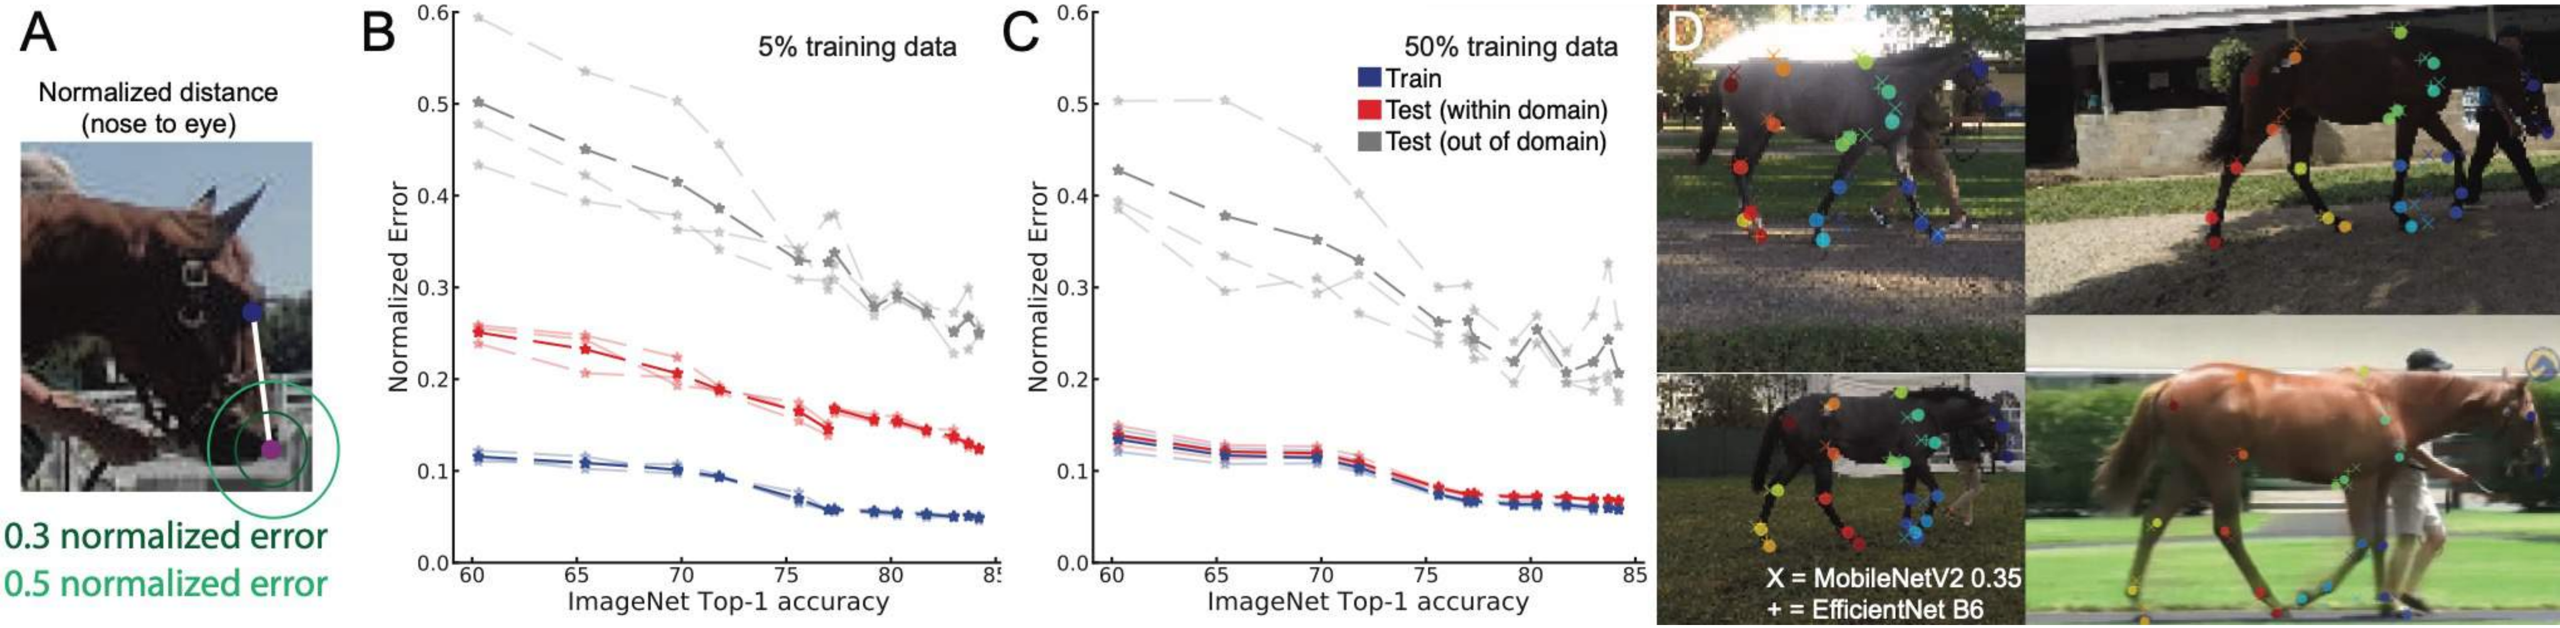

More powerful architectures generalize better

Mathis*, Biasi* et al. 2021 WACV

Transfer learning (using pretrained ImageNet models), gives a 2X boost on out-of-domain data vs. from scratch training

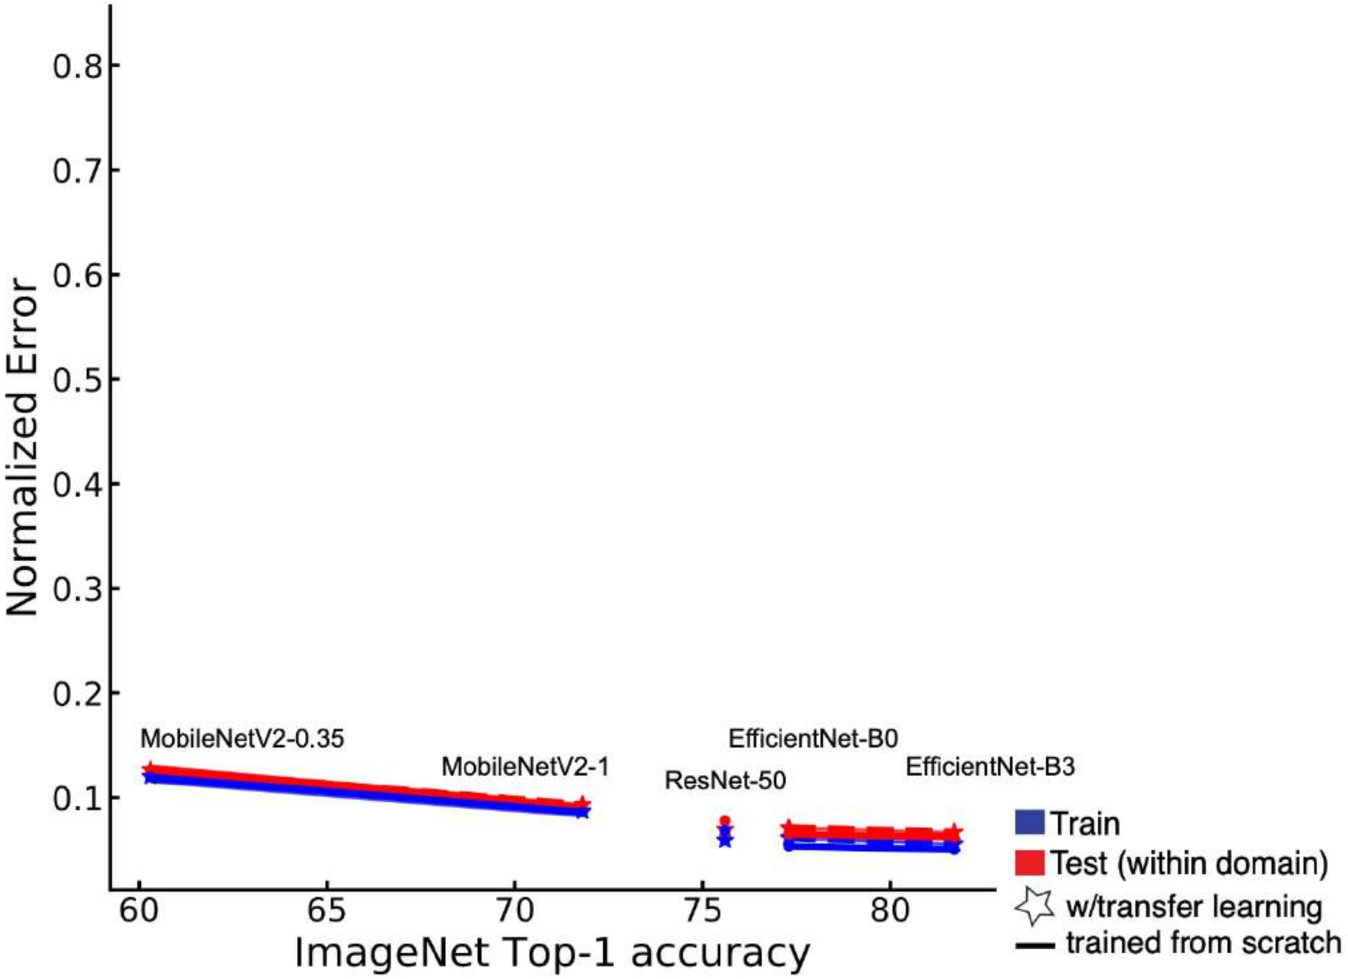

Mathis*, Biasi* et al. 2021 WACV

Using transfer learning boosts out of domain robustness
(bonus: 6X shorter training times, and less data required)

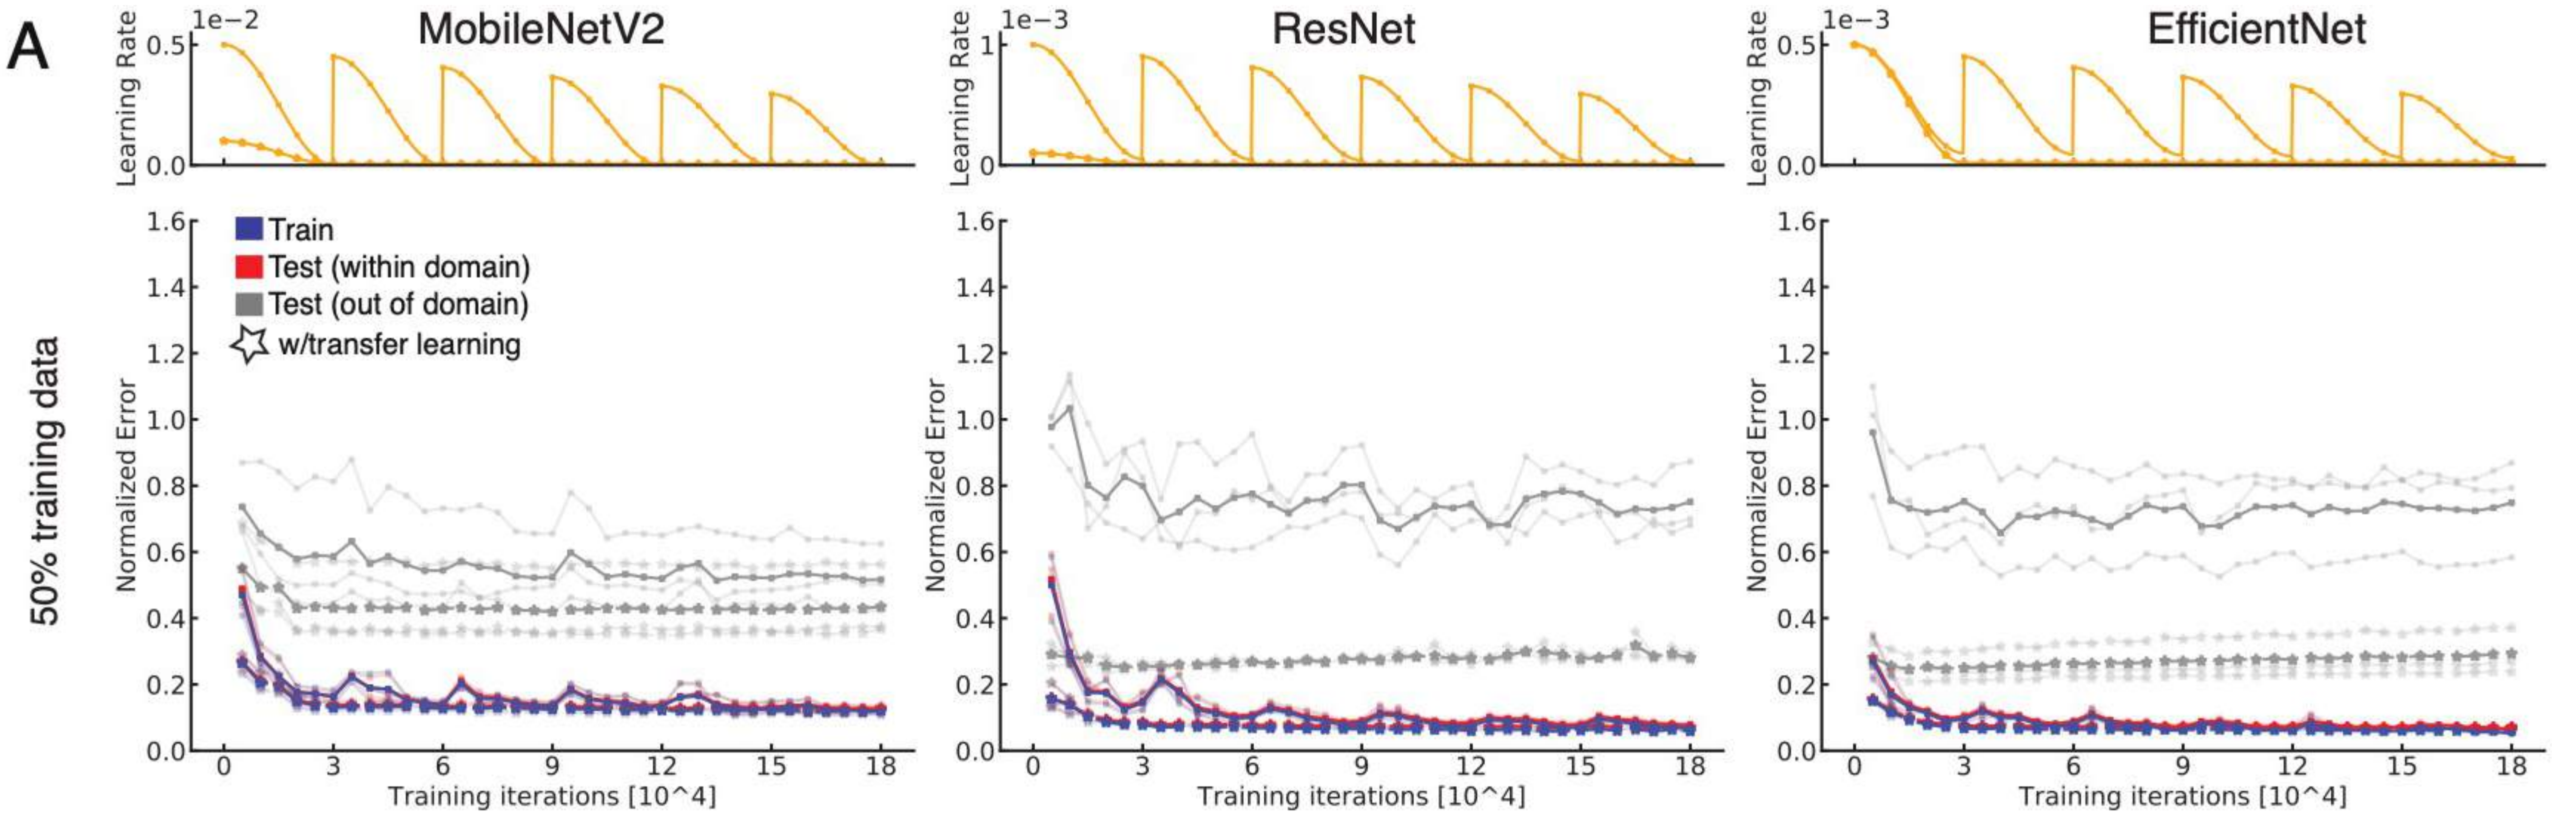

Horse-C: an animal pose estimation corruption benchmark for robustness

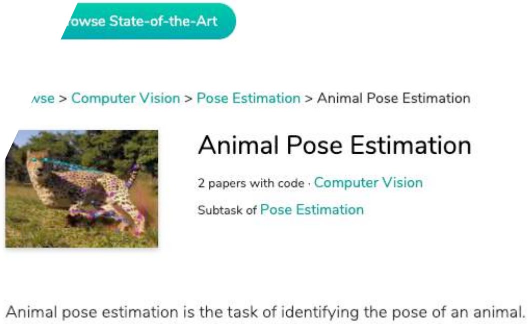

we present Horse-C to contrast the domain shift inherent in the Horse-10 dataset with domain shift induced by common image corruptions

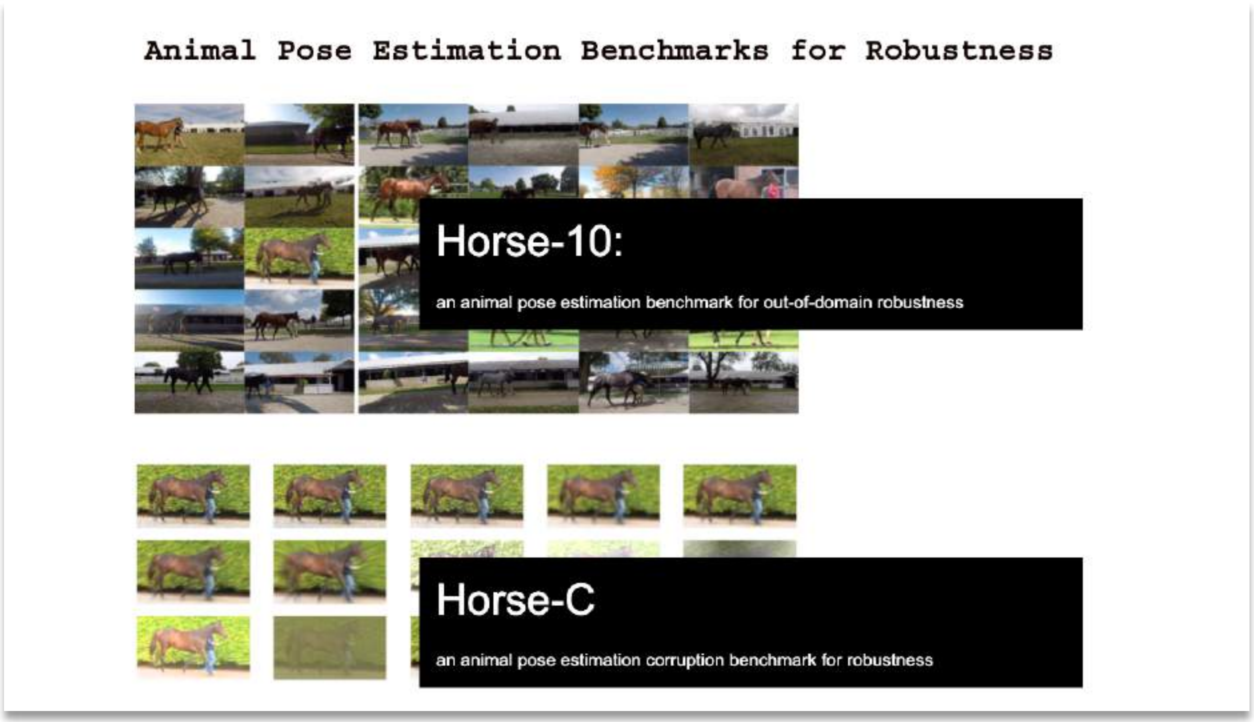

Paper, data, etc. at: http://horse10.deeplabcut.org/

~2 min talk by first author, Tom Biasi:
https://www.youtube.com/watch?v=pM6Z-ASiI2Y

Continual learning; adding small amounts of data from a new source to bring it "within domain"

Trained on this video only....

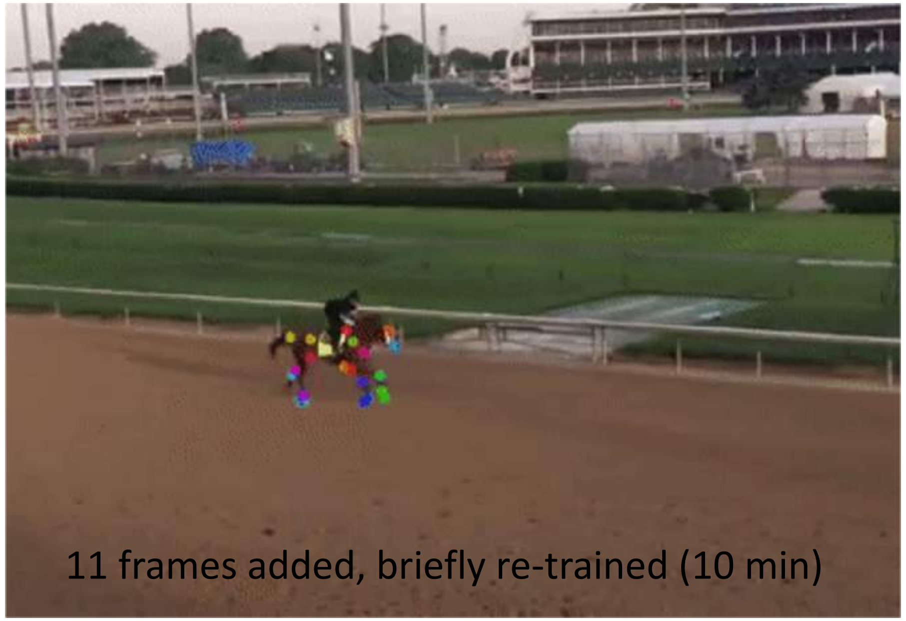

11 frames added, briefly re-trained (10 min)

Transfer learning: a more robust & generalizable network, plus less data needed

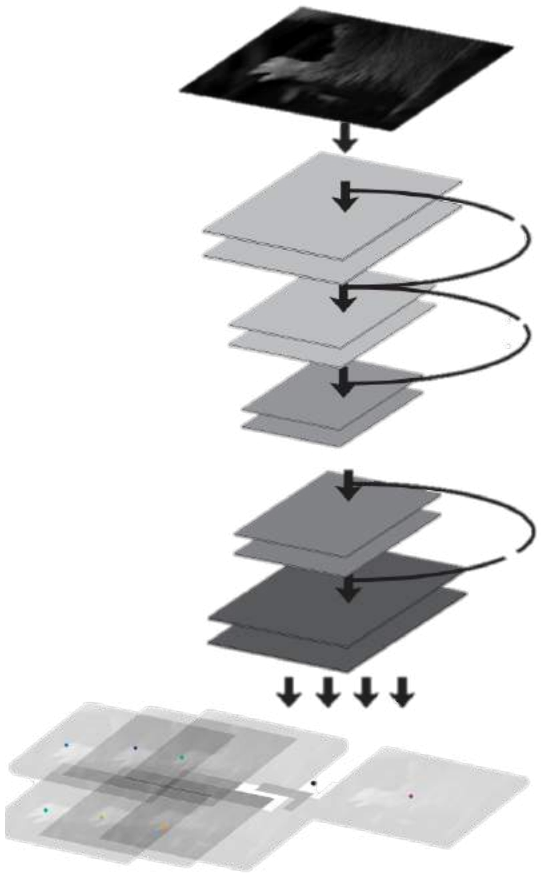

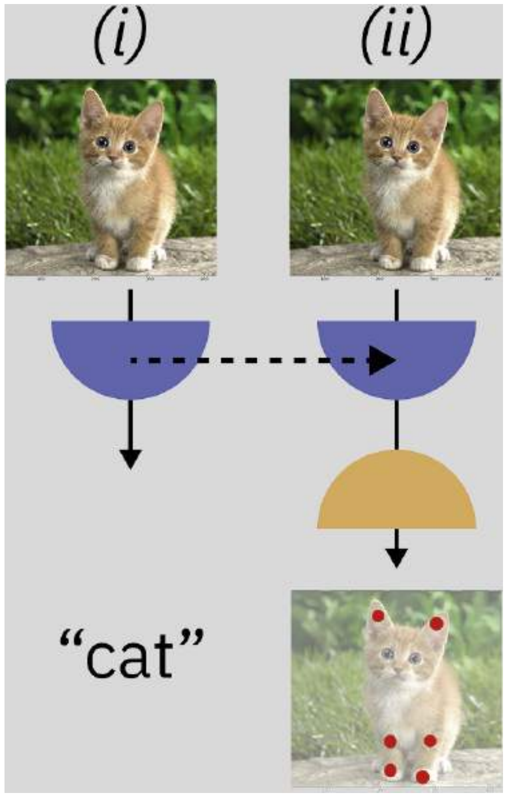

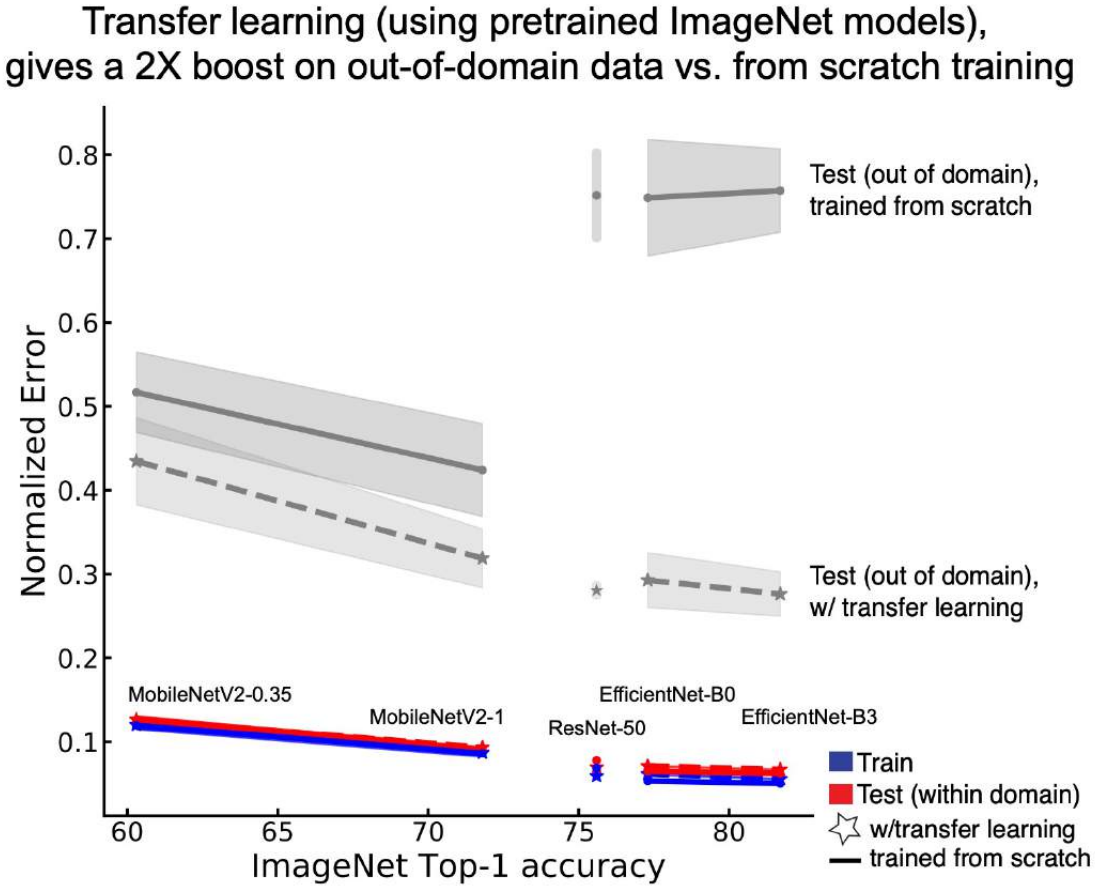

Mathis et al. 2020 IMCL workshop
Mathis, Biasi et al WACV 2021

Google Colab demo!

https://colab.research.google.com/

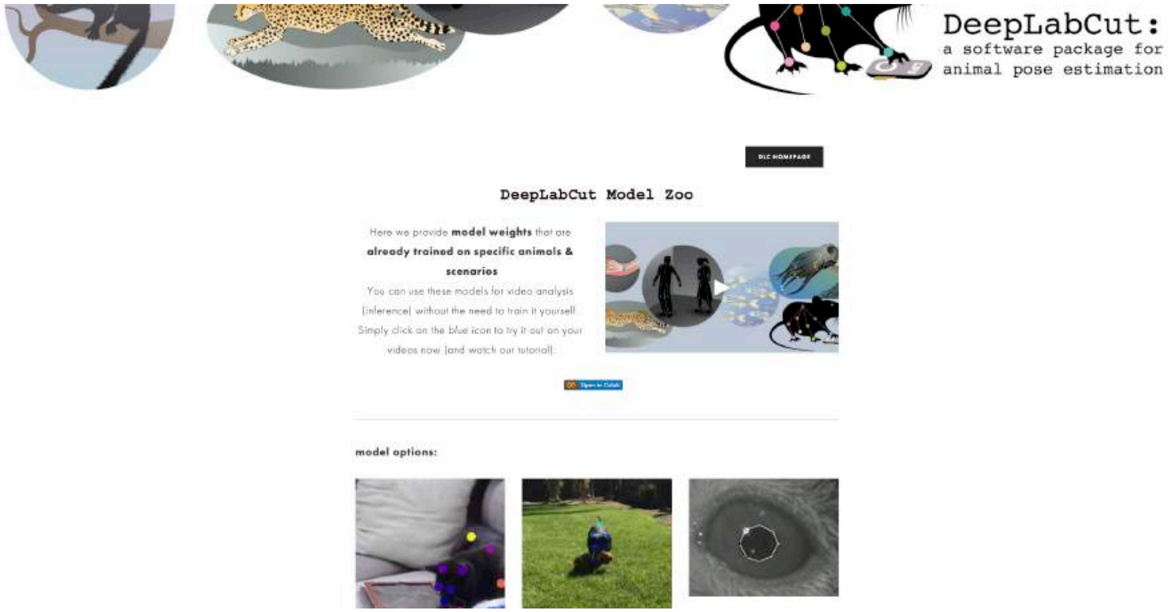

http://modelzoo.deeplabcut.org