

What Can You Learn from Your Muscles? Learning Visual Representation from Human Interactions

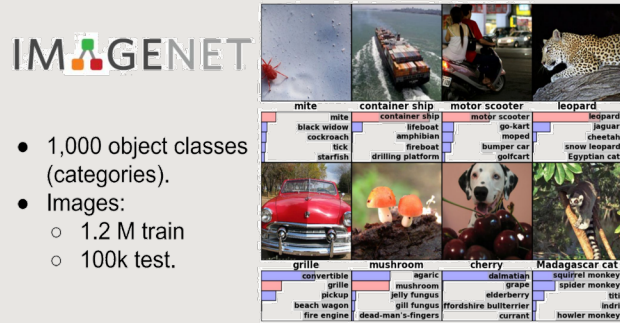
Kiana Ehsani^{*,#}, Daniel Gordon^{*}, Thomas Nguyen^{*}, Roozbeh Mottaghi^{*,#}, Ali Farhadi^{*}

^{*}University of Washington, [#]Allen Institute for Artificial Intelligence

ICLR 2021


Visual Representation

ImageNet Challenge




- 1,000 object classes (categories).
- Images:
 - 1.2 M train
 - 100k test.


He, Kaiming, et al. "Deep residual learning for image recognition." *CVPR*'16



Noroozi, et al. "Unsupervised learning of visual representations by solving jigsaw puzzles." *ECCV*'16.



Zhang, Richard, et al. "Colorful image colorization." *ECCV*'16.



He, et al. "Momentum contrast for unsupervised visual representation learning." *CVPR*'20.

Oord, et al. "Representation learning with contrastive predictive coding." *Arxiv*'18.

Wu, et al. "Unsupervised feature learning via non-parametric instance discrimination." *CVPR*'18.

Learning From the Muscles

77 hours

60 fps

35 participants

Variety of unlabeled actions

Cooking



Playing
Pool

Biking

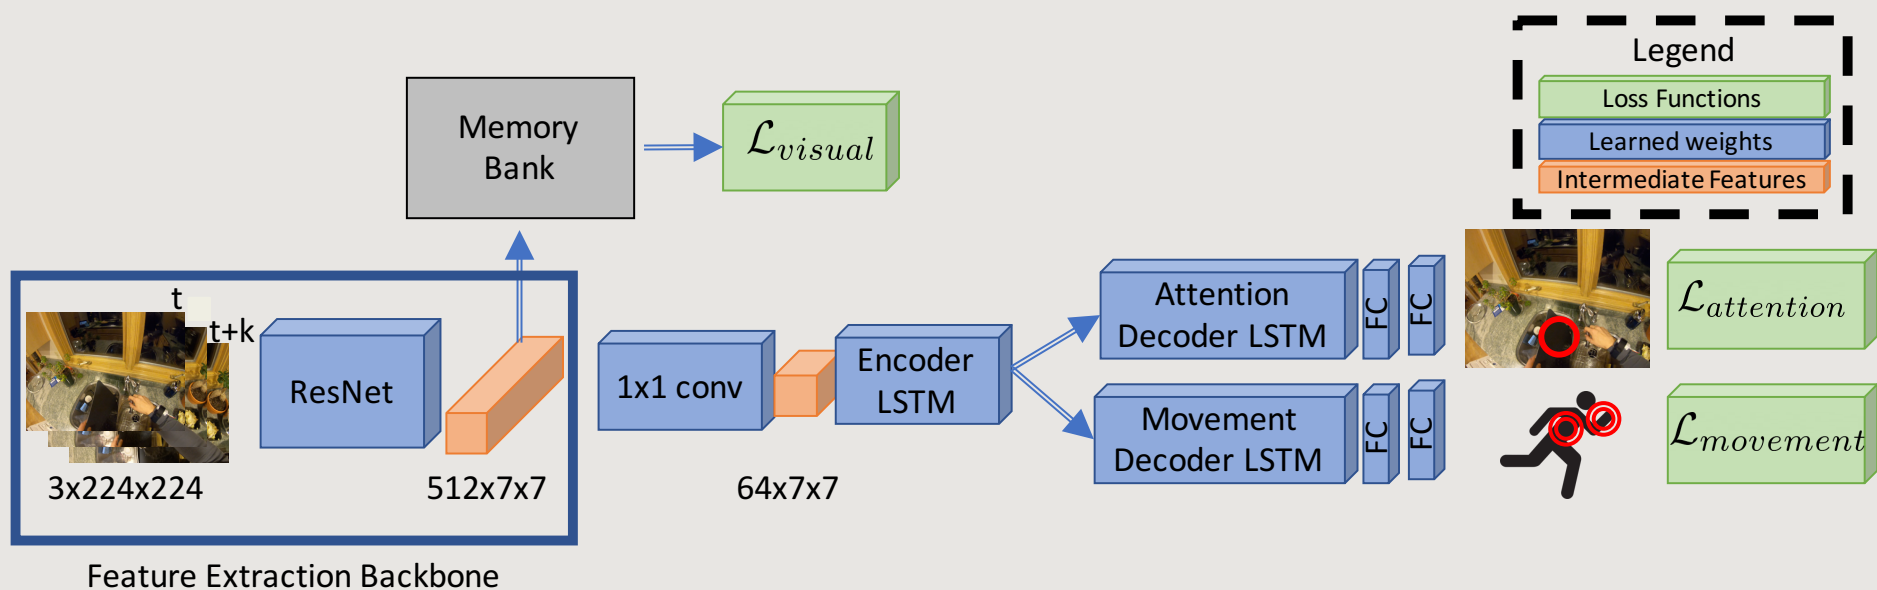
Shoveling
Snow


Climbing

And Many
More




Model






Ground-Truth

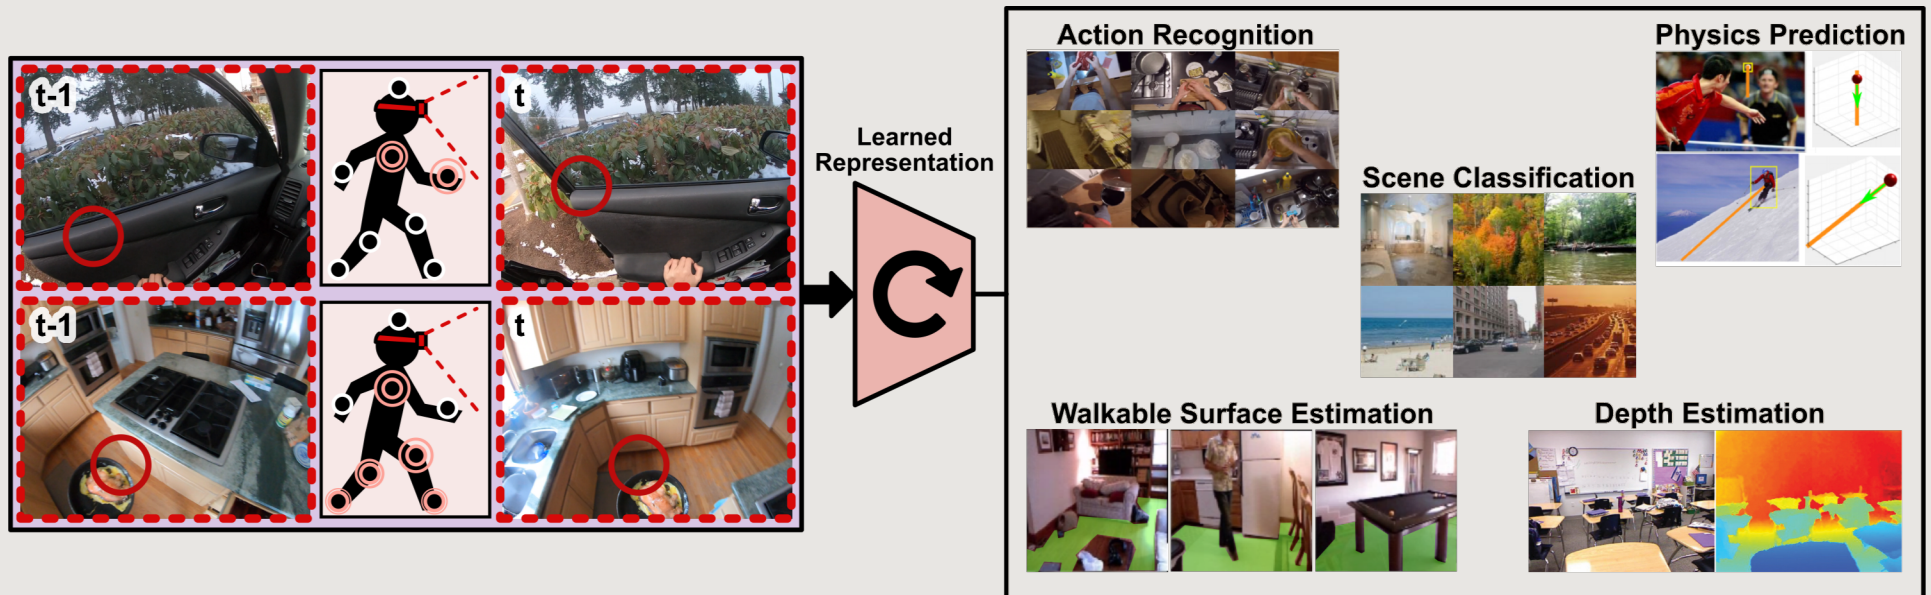


Predictions


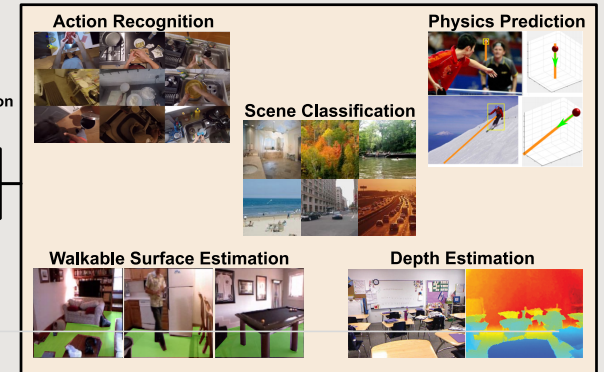


*CIRCLES ON THE HUMAN FIGURE SHOW MOVEMENT IN THE PARTS.
EACH VIDEO IS PLAYED THREE TIMES.*

Results



Results



[MoCo] He, Kaiming, et al. "Momentum contrast for unsupervised visual representation learning." CVPR'20.

The numbers for depth estimation are errors and for easier interpretation we are plotting the inverse of the error.

What Can You Learn from Your Muscles?

Learning Visual Representation from Human Interactions

- Thank you for listening to our presentation!
- Our code and dataset are released:

<https://github.com/ehsanik/muscleTorch>

