




PRESENTER: **Yifeng Zhu**

VIOLA uses **object proposal priors** to learn a **closed-loop visuomotor policy** to **make coffee** with only **50 raw demonstrations**



Motivation

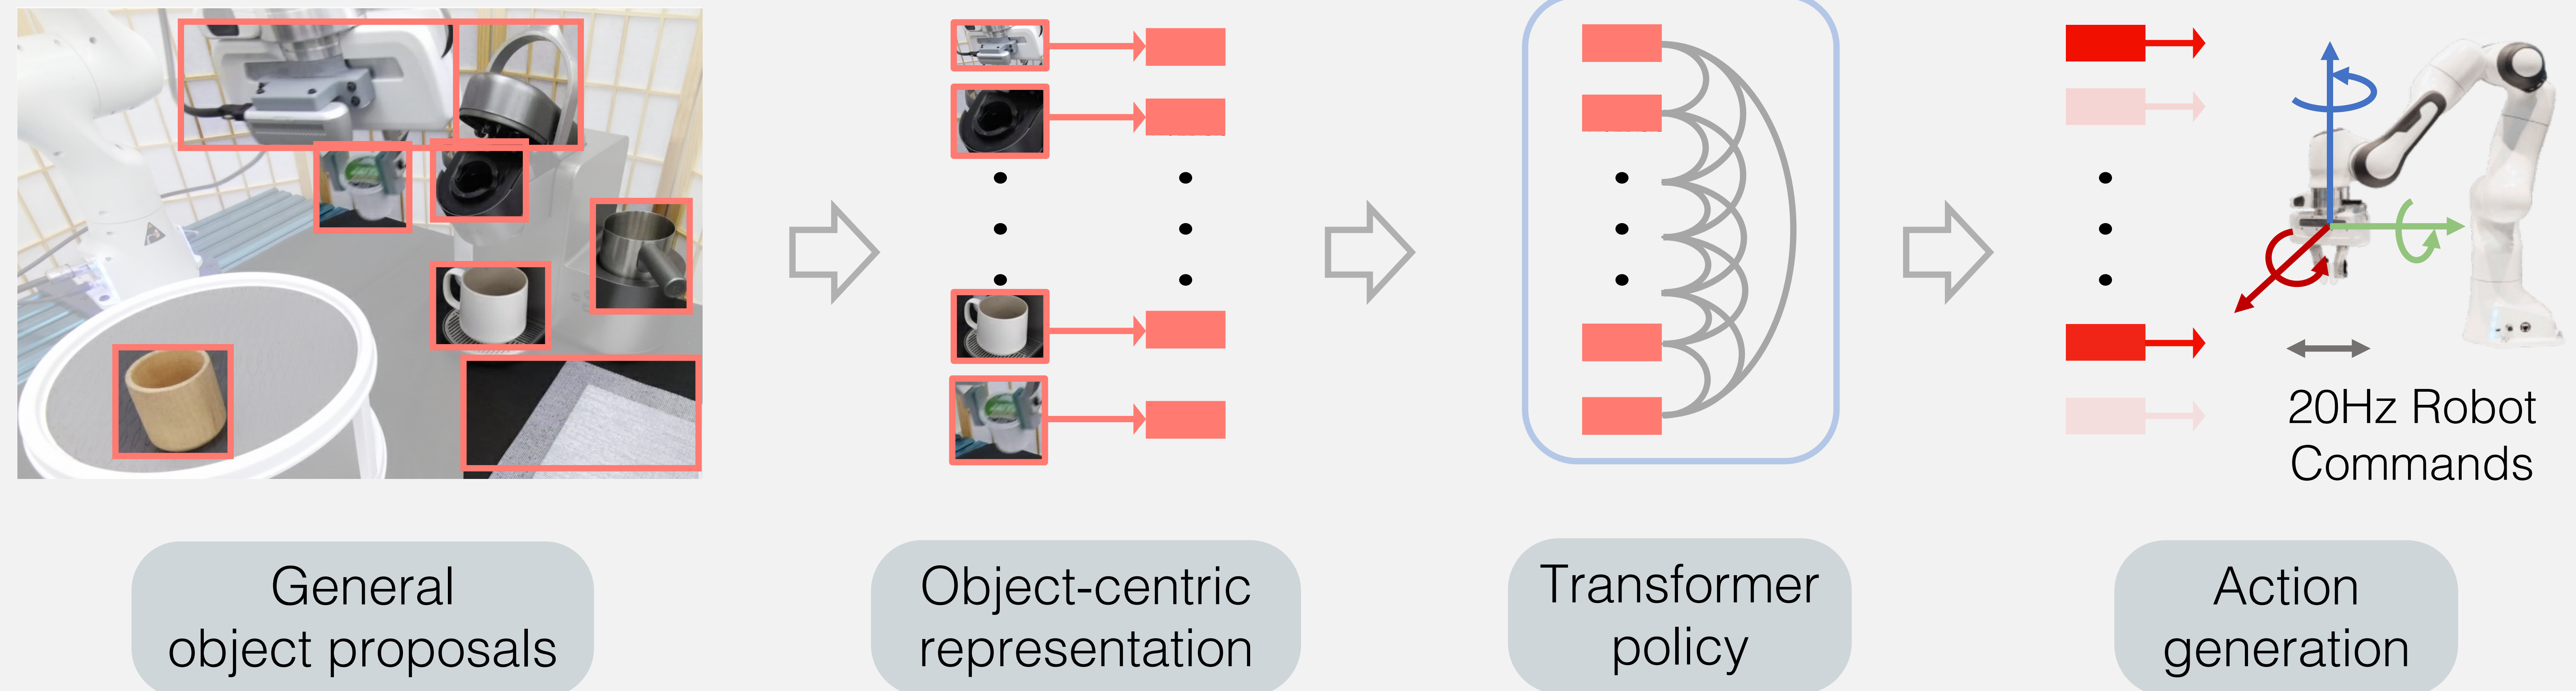
- Imitate closed-loop visuomotor policies efficiently for manipulation
- Reduce spurious correlation in visuomotor learning
- Improve policy generalization from a handful of demonstrations

Insights

- Object-centric priors facilitate more efficient and robust policy inference
- Pre-trained RPN from large-scale image datasets captures general object priors
- Transformer self-attention mechanism selects task-relevant object features

VIOLA: Imitation Learning for Vision-Based Manipulation with Object Proposal Priors

Yifeng Zhu*, Abhishek Joshi, Peter Stone, Yuke Zhu



Experiments



Dining-PlateFork
 BC-RNN: 36.7
 VIOLA: 76.7



Dining-Bowl
 BC-RNN: 20.0
 VIOLA: 60.0



Make-Coffee
 BC-RNN: 0.0
 VIOLA: 60.0

Success rates (%)

■ BC-RNN ■ VIOLA

Laundry List

- Impedance controllers for behavior cloning
- Suboptimal data actually helps to learn close-loop behaviors
- Generalization is limited to same task setup from training

VIOLA Model

