


Error-Aware Imitation Learning from Teleoperation Data for Mobile Manipulation

Josiah Wong Albert Tung Andrey Kurenkov Ajay Mandlekar Li Fei-Fei Silvio Savarese Roberto Martín-Martín

Website: sites.google.com/view/il-for-mm/home




What's the Problem?

Domain



Mobile Manipulation (MM)

- + endows robots with diverse capabilities
- suffers from vast state space difficult to explore



Imitation Learning


- + promising in static manipulation and navigation
- unexplored extensively in MM



Challenges

Intuitive manipulative and locomotive control is **non-trivial**^[1]

Possibly **significant covariate shift** at test time^[2]



What's our Insight?

(1) Designing a **novel teleoperation interface (Mobile Manipulation RoboTurk)** enables us to...

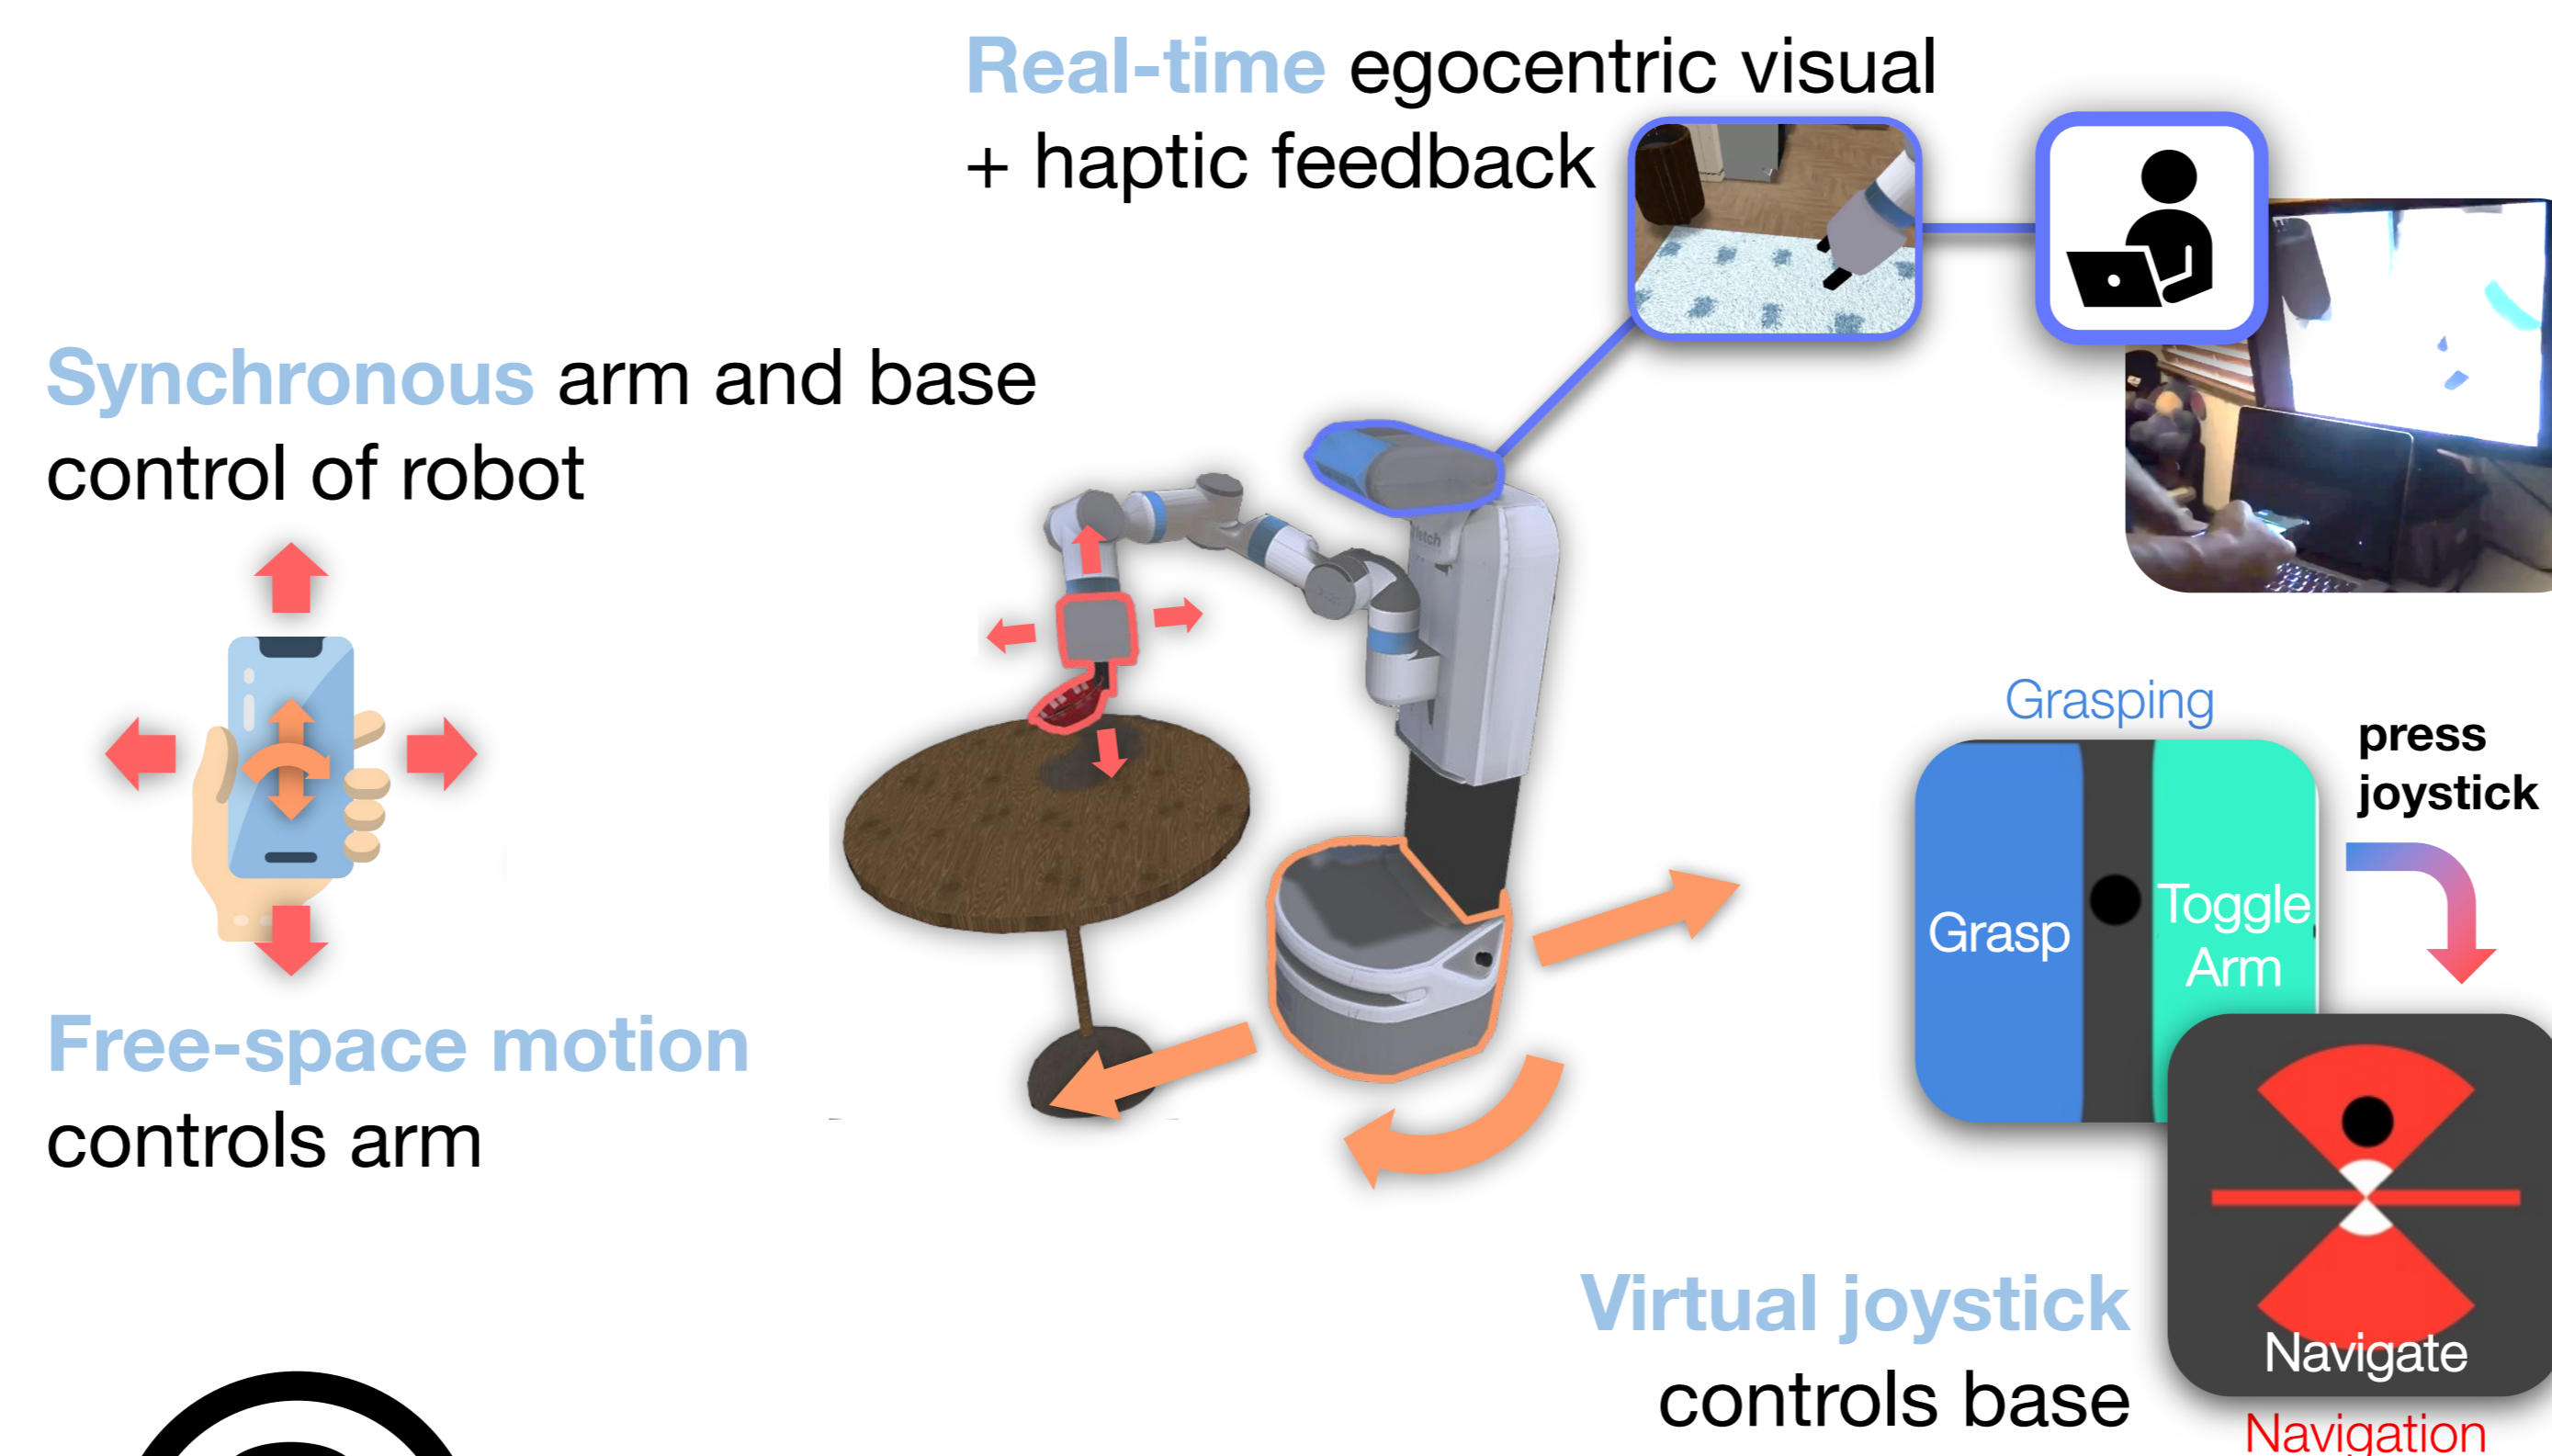
(2) Collect a **first-of-its-kind large-scale demonstration dataset** that can be used to...

(3) train **performant, error-aware imitation learning policies**



Mobile Manipulation RoboTurk (MoMaRT)

MoMaRT enables allows users to **directly teleoperate** a mobile manipulator arm + base with their smartphone^[3]



Large-Scale Dataset

We collect a **novel continuous control dataset** to accelerate future MM research

Volume
1200+ Demos
11+ Hours

Tasks

5 Long-Horizon, Multi-Stage Simulated Kitchen Tasks

Demonstrators

1 Expert
2 Suboptimal

Demo Subsets

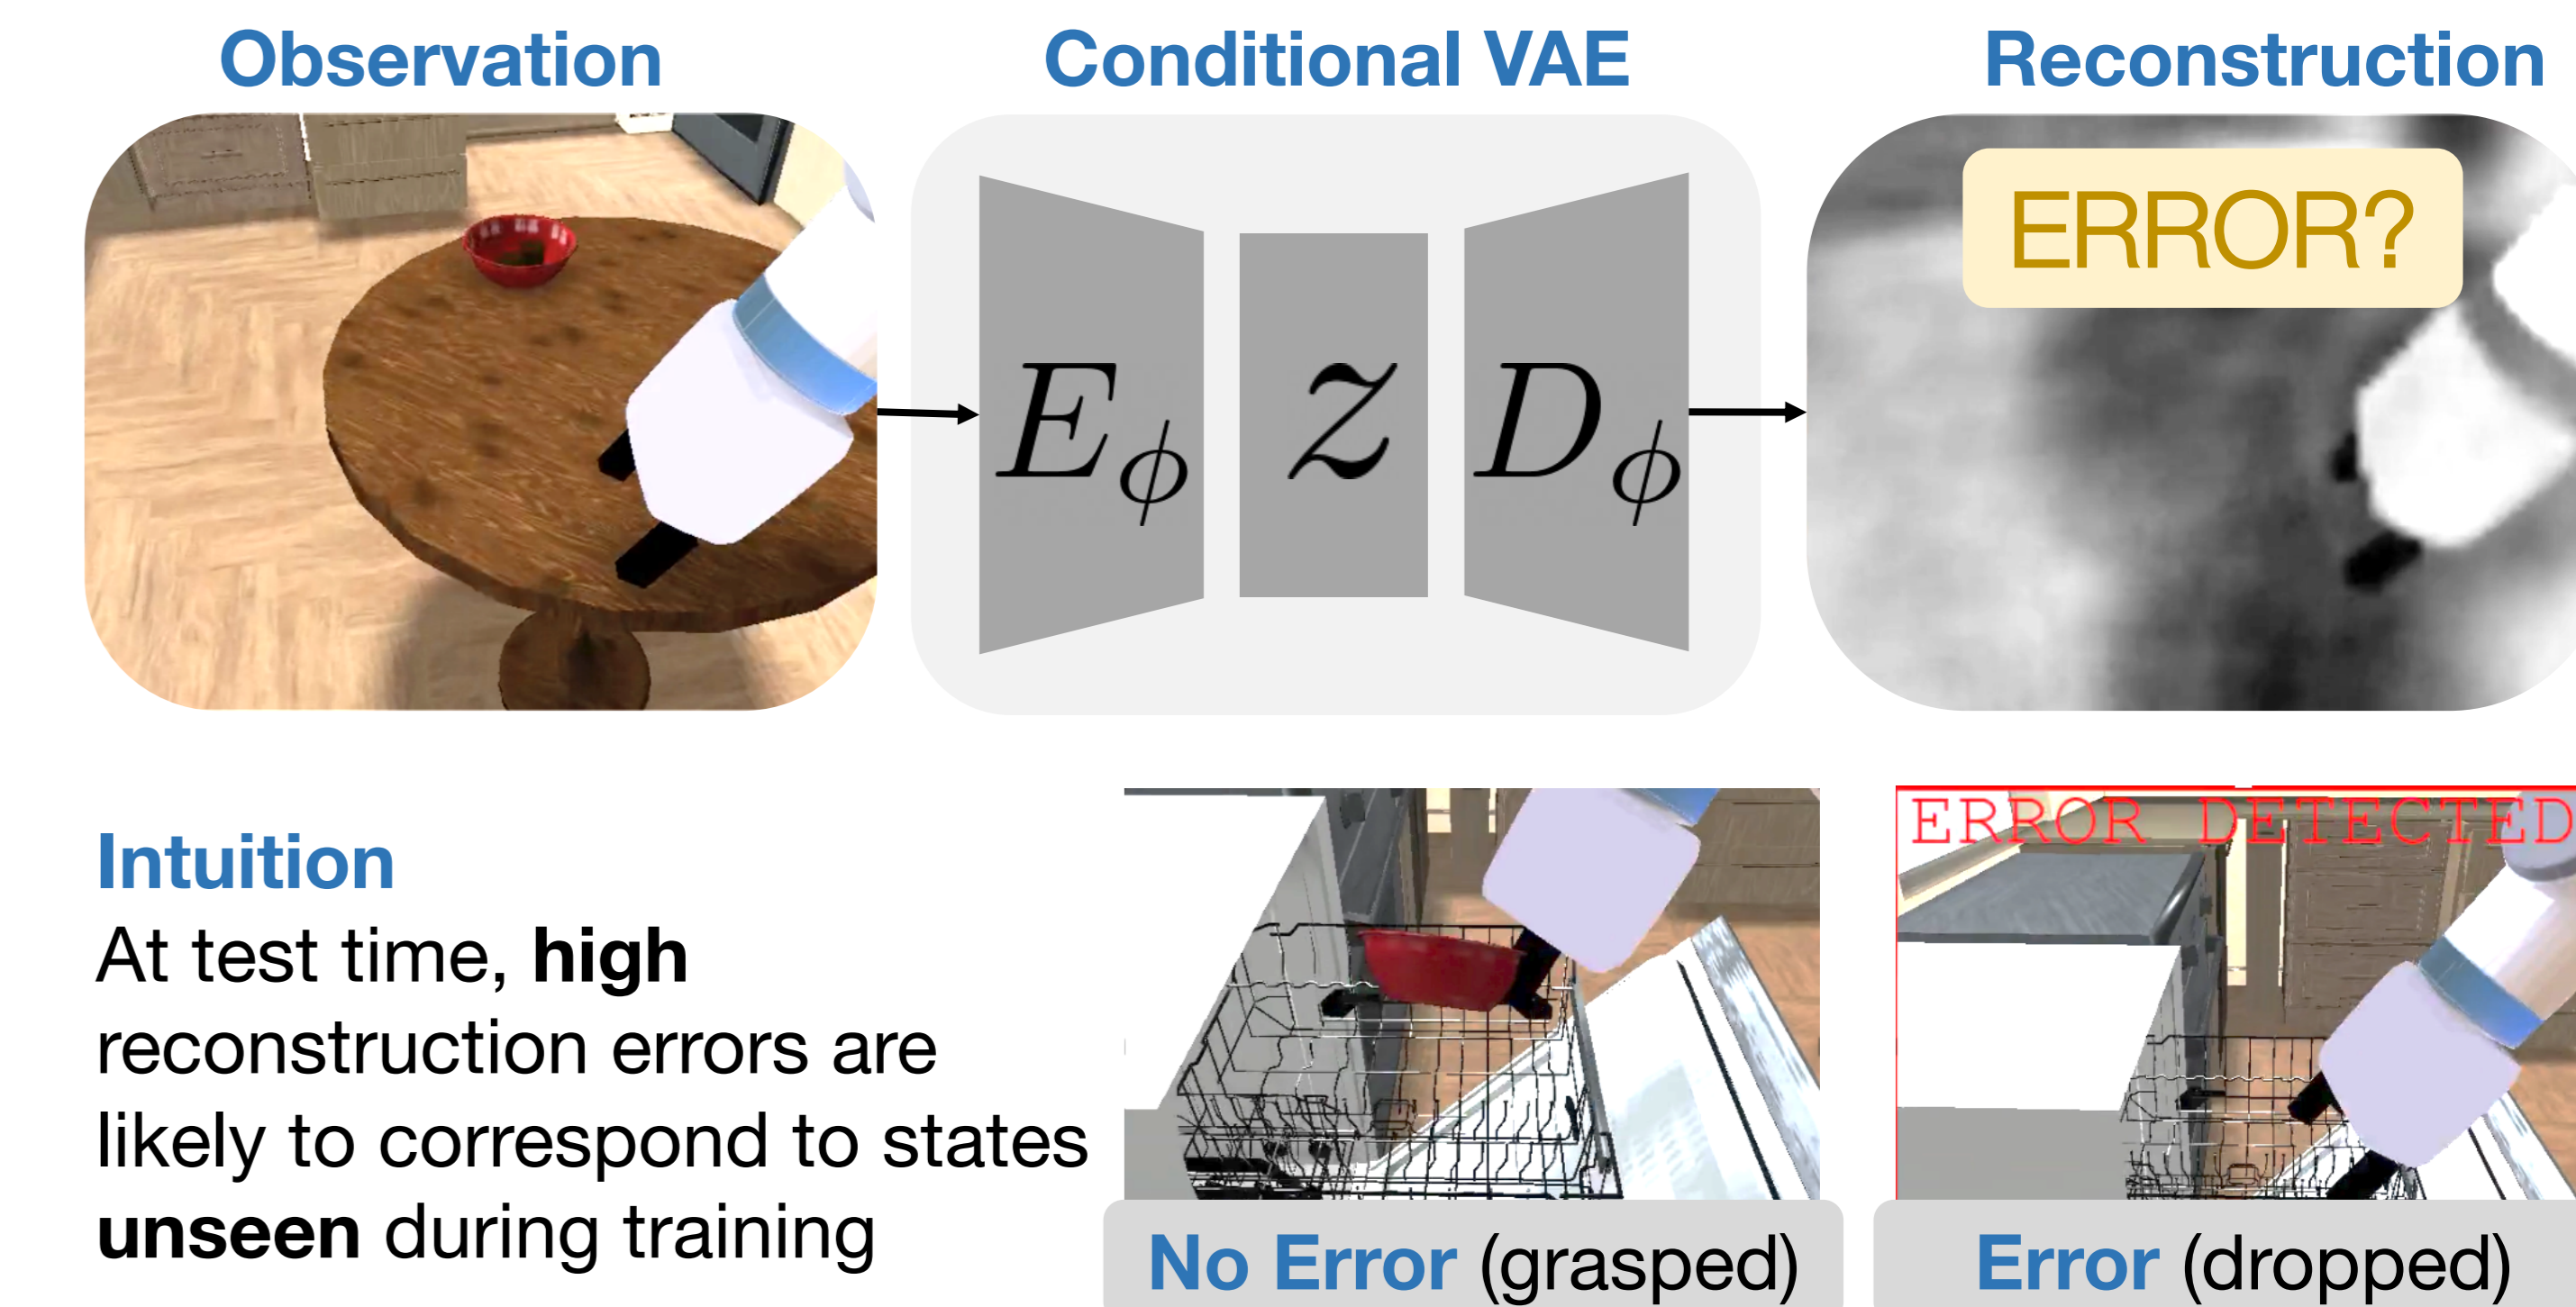
550+ Expert
550+ Suboptimal
100+ Generalization

1 Environment, Many Subtasks



Error-Aware Policies

We train a simple, but effective **error-detector** fully offline to detect out-of-distribution states^[4]



Results



48-68%
Policy Success



86-100%
Error Detection



80-100%
Error Detection
(Few-Shot Generalization)



Key Takeaways

- MoMaRT can quickly **enable large-scale MM data collection**
- Large-Scale Dataset** is publicly available to **encourage diverse MM research** (compositionality, transfer learning, ...)
- Error Detector** method can **improve MM safety** without requiring any additional online data

References

[1] E. Ratner, B. Cohen, M. Phillips, and M. Likhachev. A web-based infrastructure for recording user demonstrations of mobile manipulation tasks. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015
 [2] S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 2011.
 [3] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, et al. Roboturk: A crowdsourcing platform for robotic skill learning through imitation. In *Conference on Robot Learning*, 2018.
 [4] M. A. Lee, M. Tan, Y. Zhu, and J. Bohg. Detect, reject, correct: Crossmodal compensation of corrupted sensors. *arXiv preprint arXiv:2012.00201*, 2020.