Beyond the Pixels: Exploring the Effect of Video File Corruptions on Model Robustness




Trenton Chang Daniel Y. Fu Yixuan Li Christopher Ré

Stanford University

Workshop on Adversarial Robustness in the Real World, ECCV 2020

Outline

- **Video robustness: why file corruption?**
- **Setup**
  - Model evaluation
  - Simulating file corruptions
- **Results**
  - Effect of file corruption on model performance
  - Qualitative analysis of corrupted videos
  - Quantitative analysis of corrupted videos

Video robustness: not just a pixel-space problem

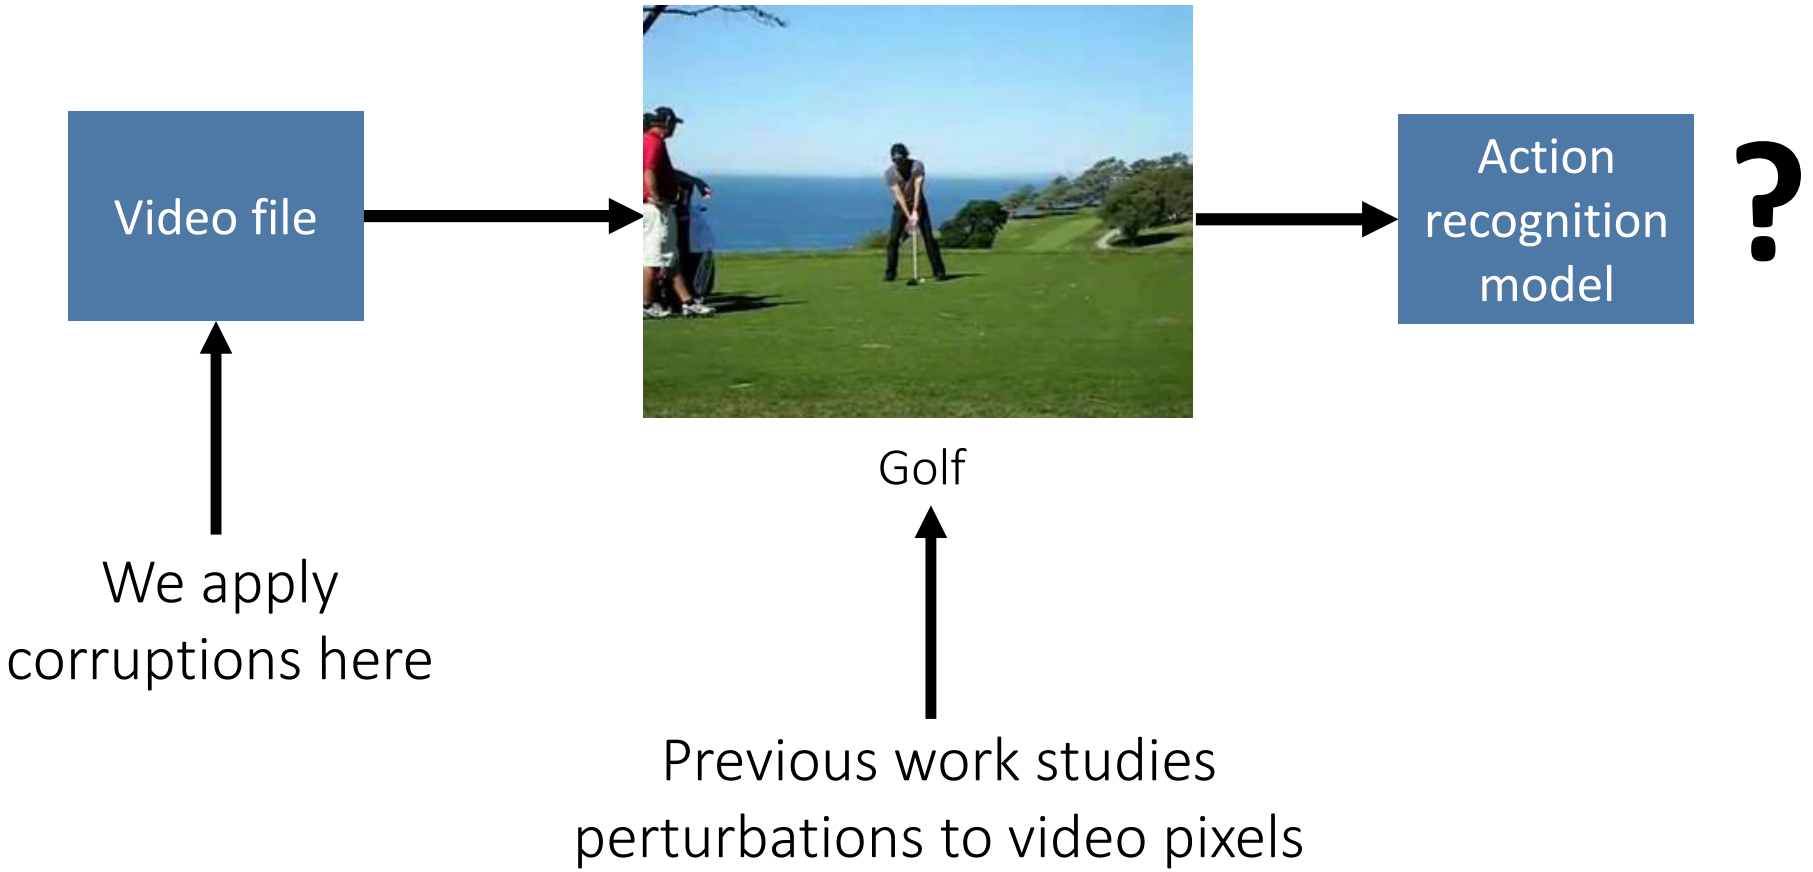

Goal: Simulate real-world file corruptions and measure their effect on video model robustness.

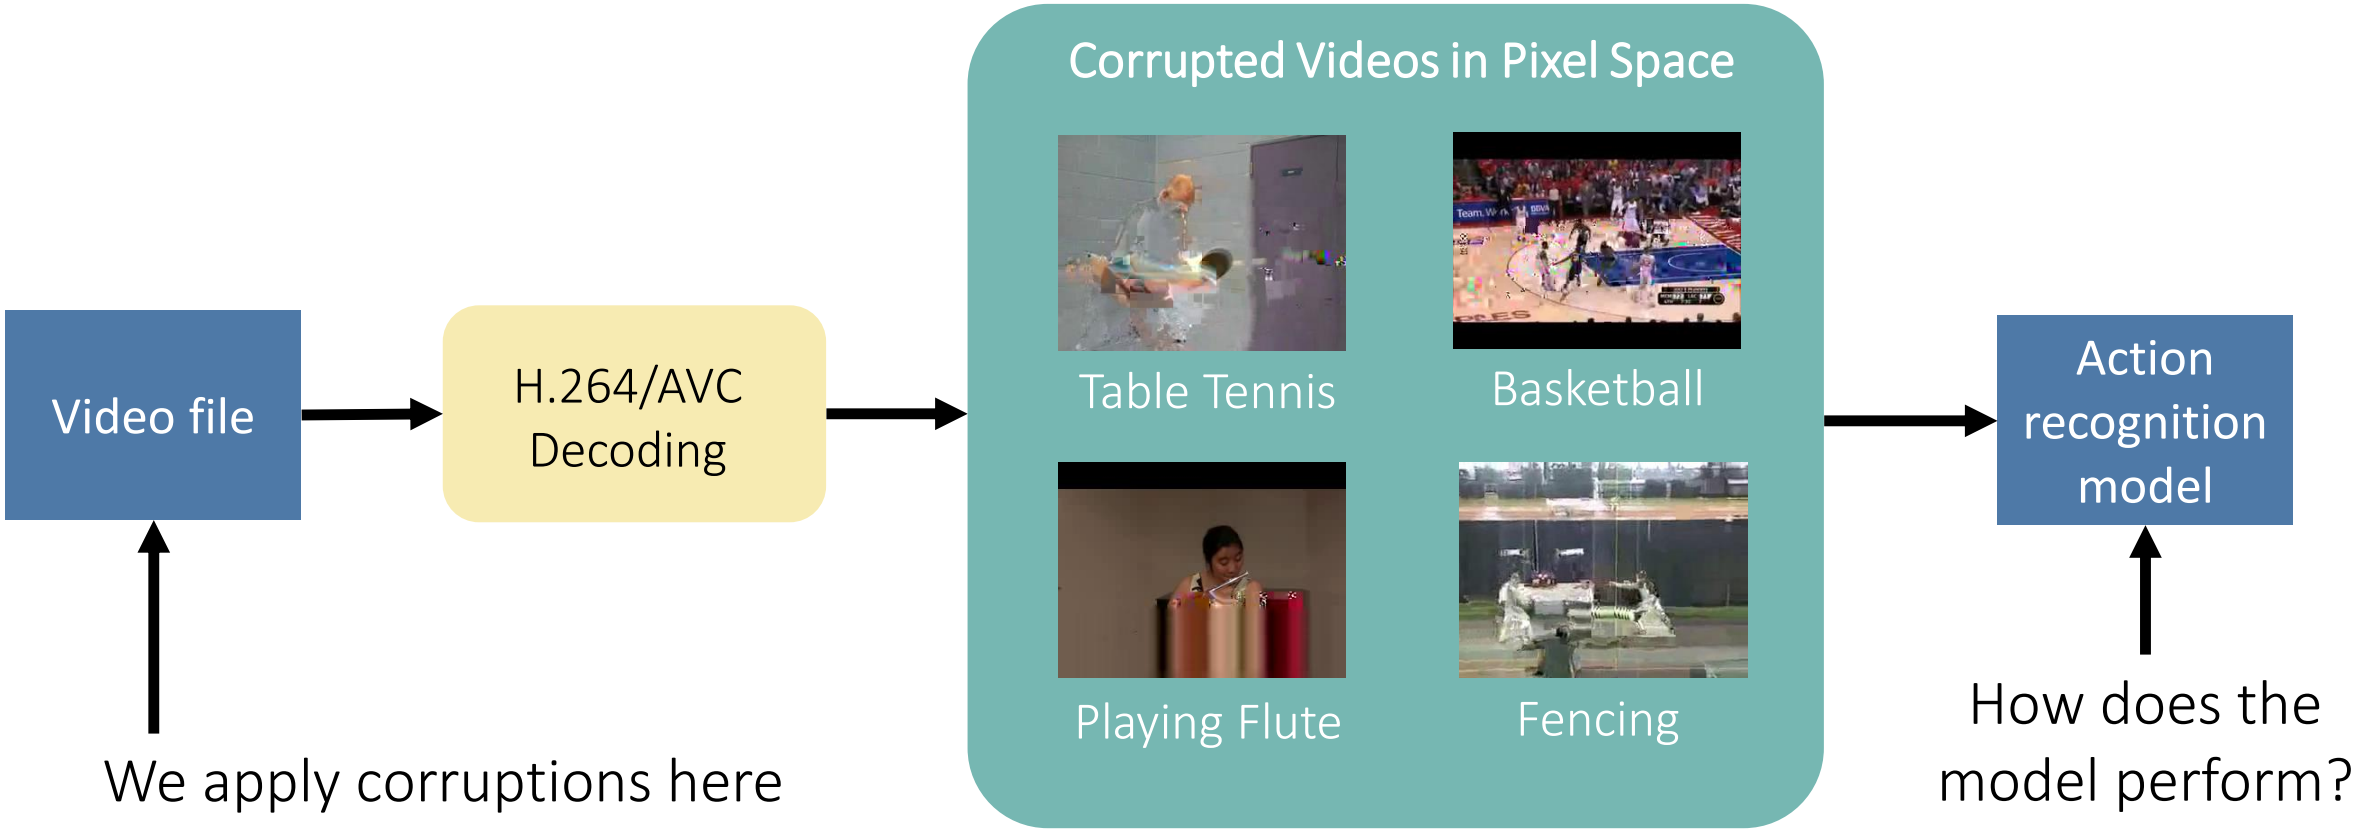

Setup

Model evaluation pipeline

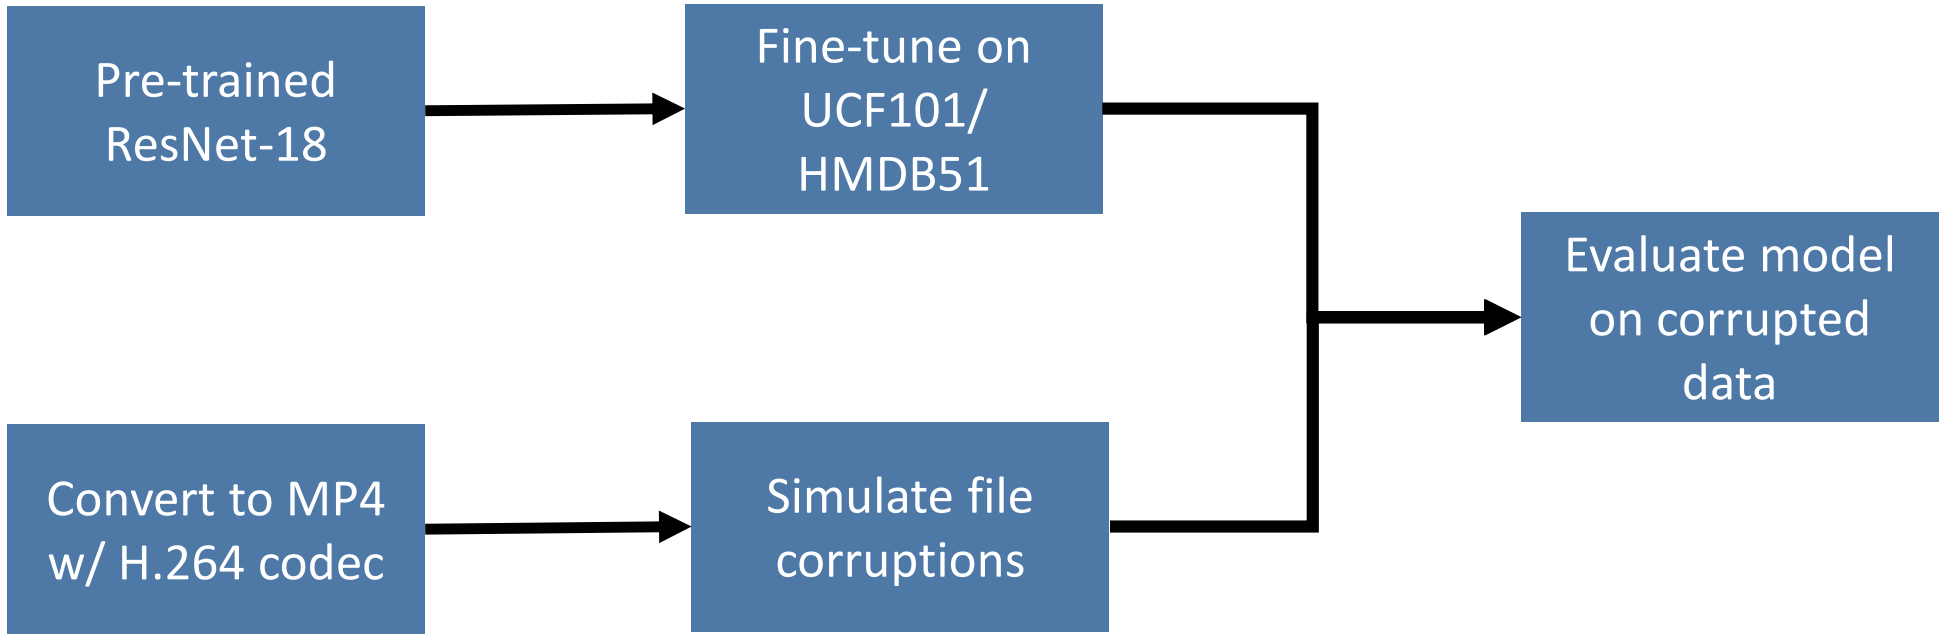

Simulating two types of file corruptions

Random corruption

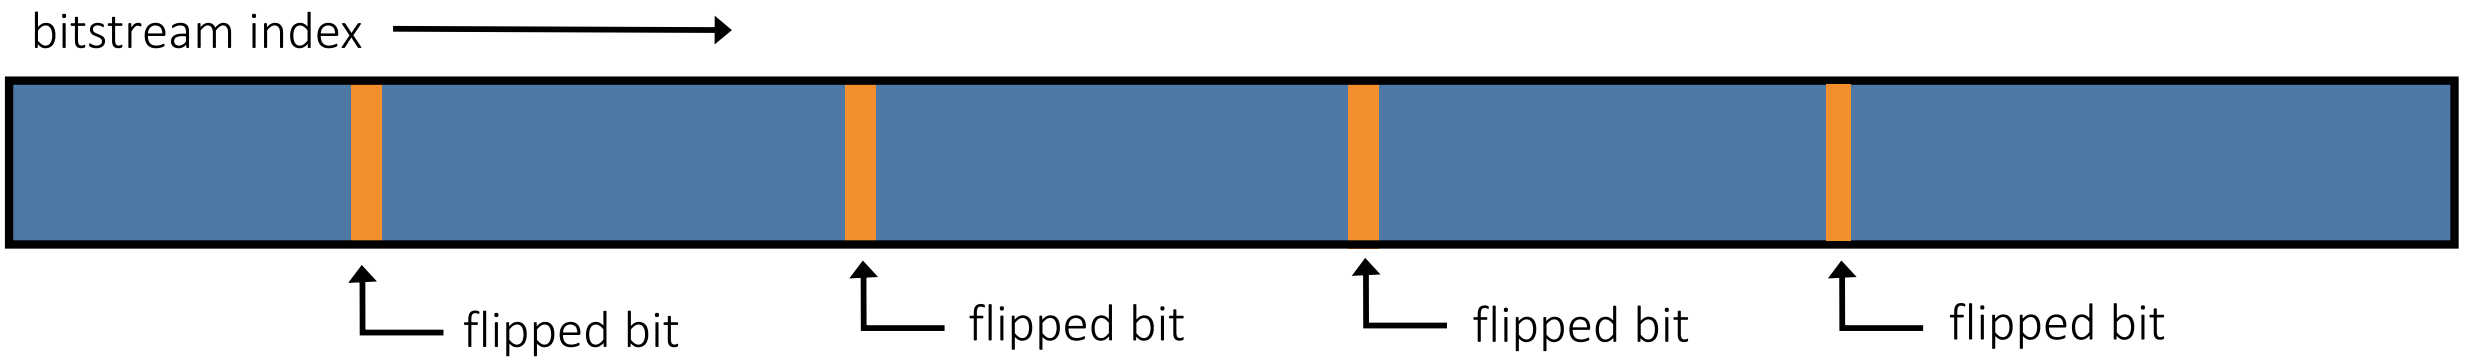

Contiguous corruption

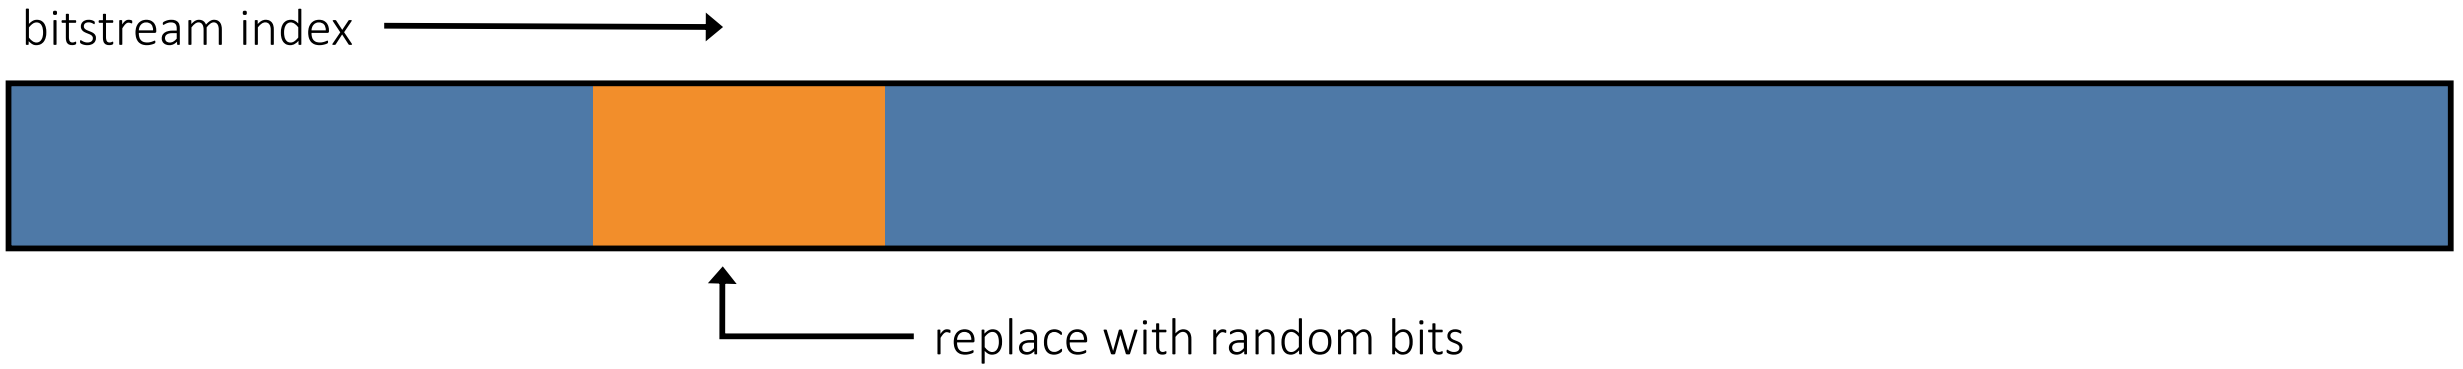

Experiments: vary total length of orange segments (corruption proportion)

Results & Discussion

Accuracy drops as corruption proportion increases

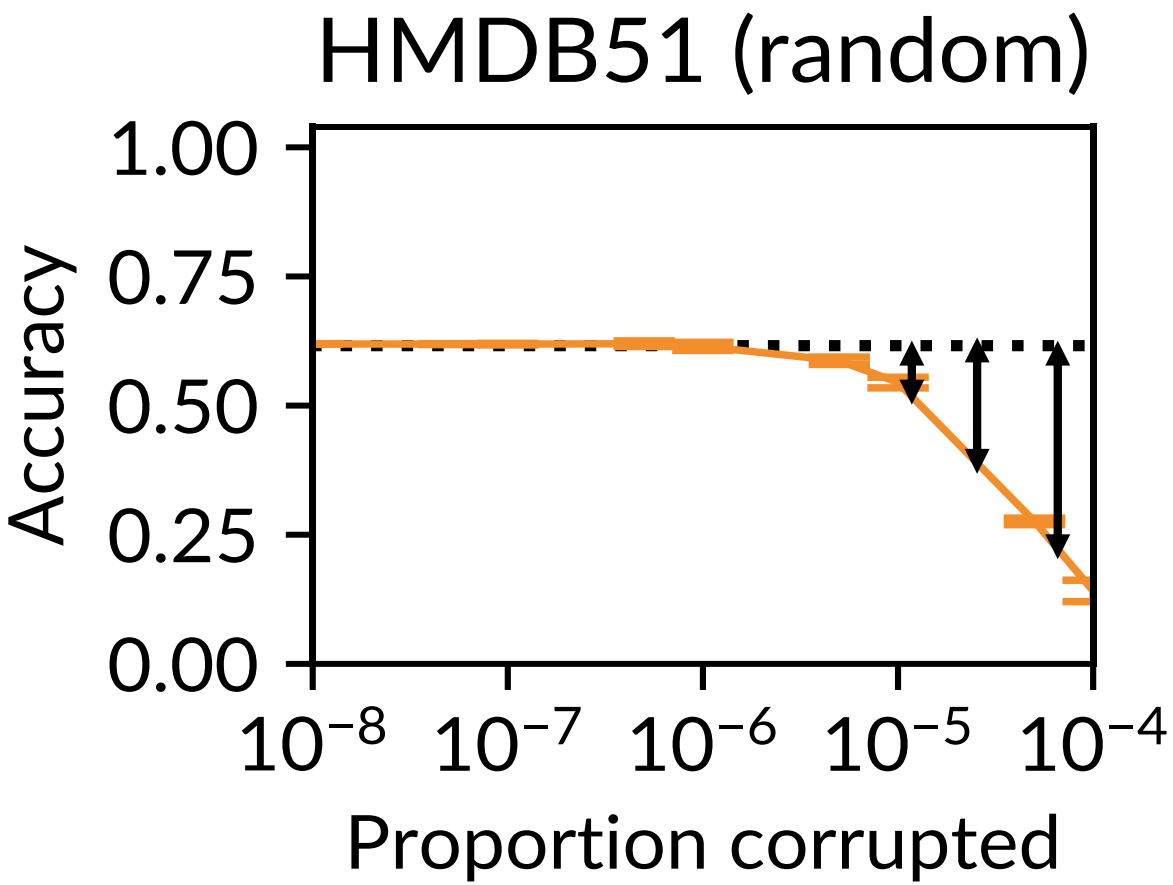

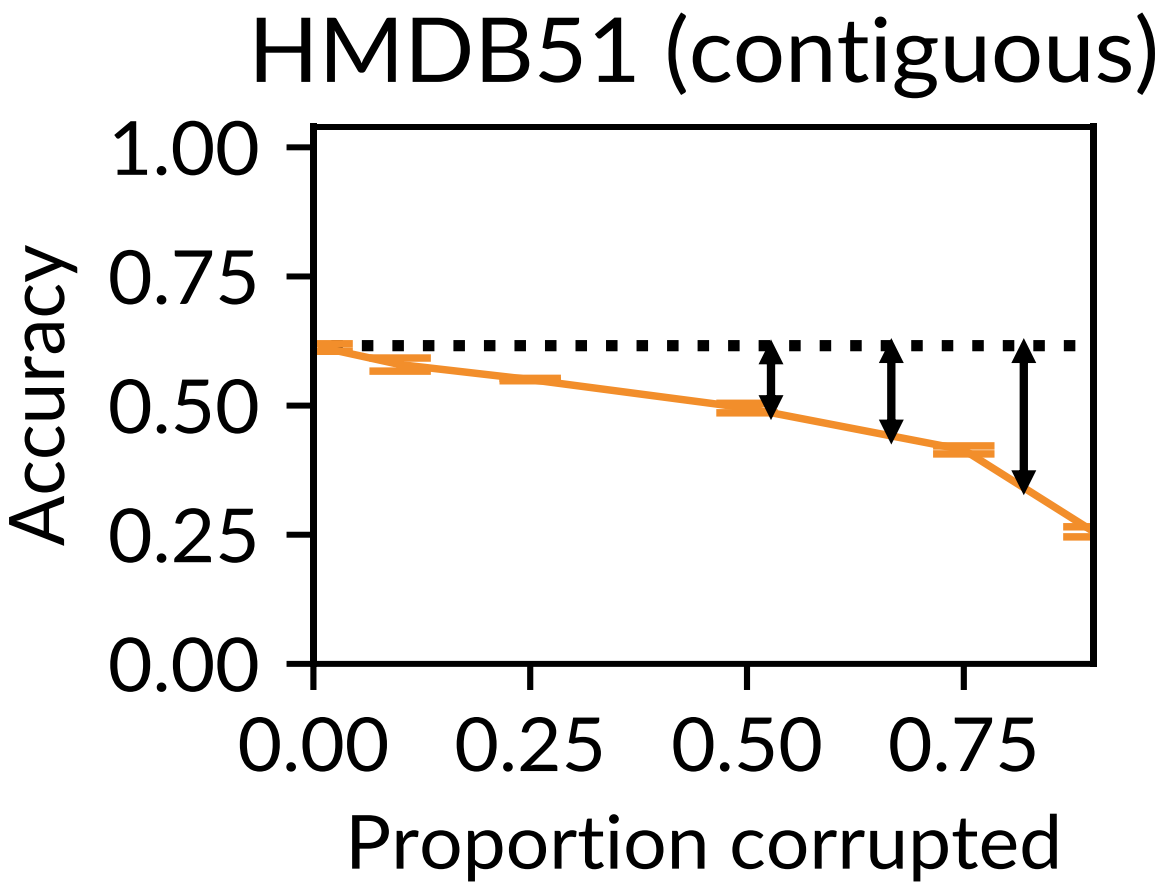

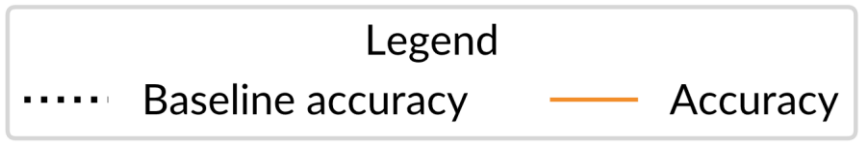

Model accuracy drops as corruption proportion increases

Model errors correlate with corruption proportion

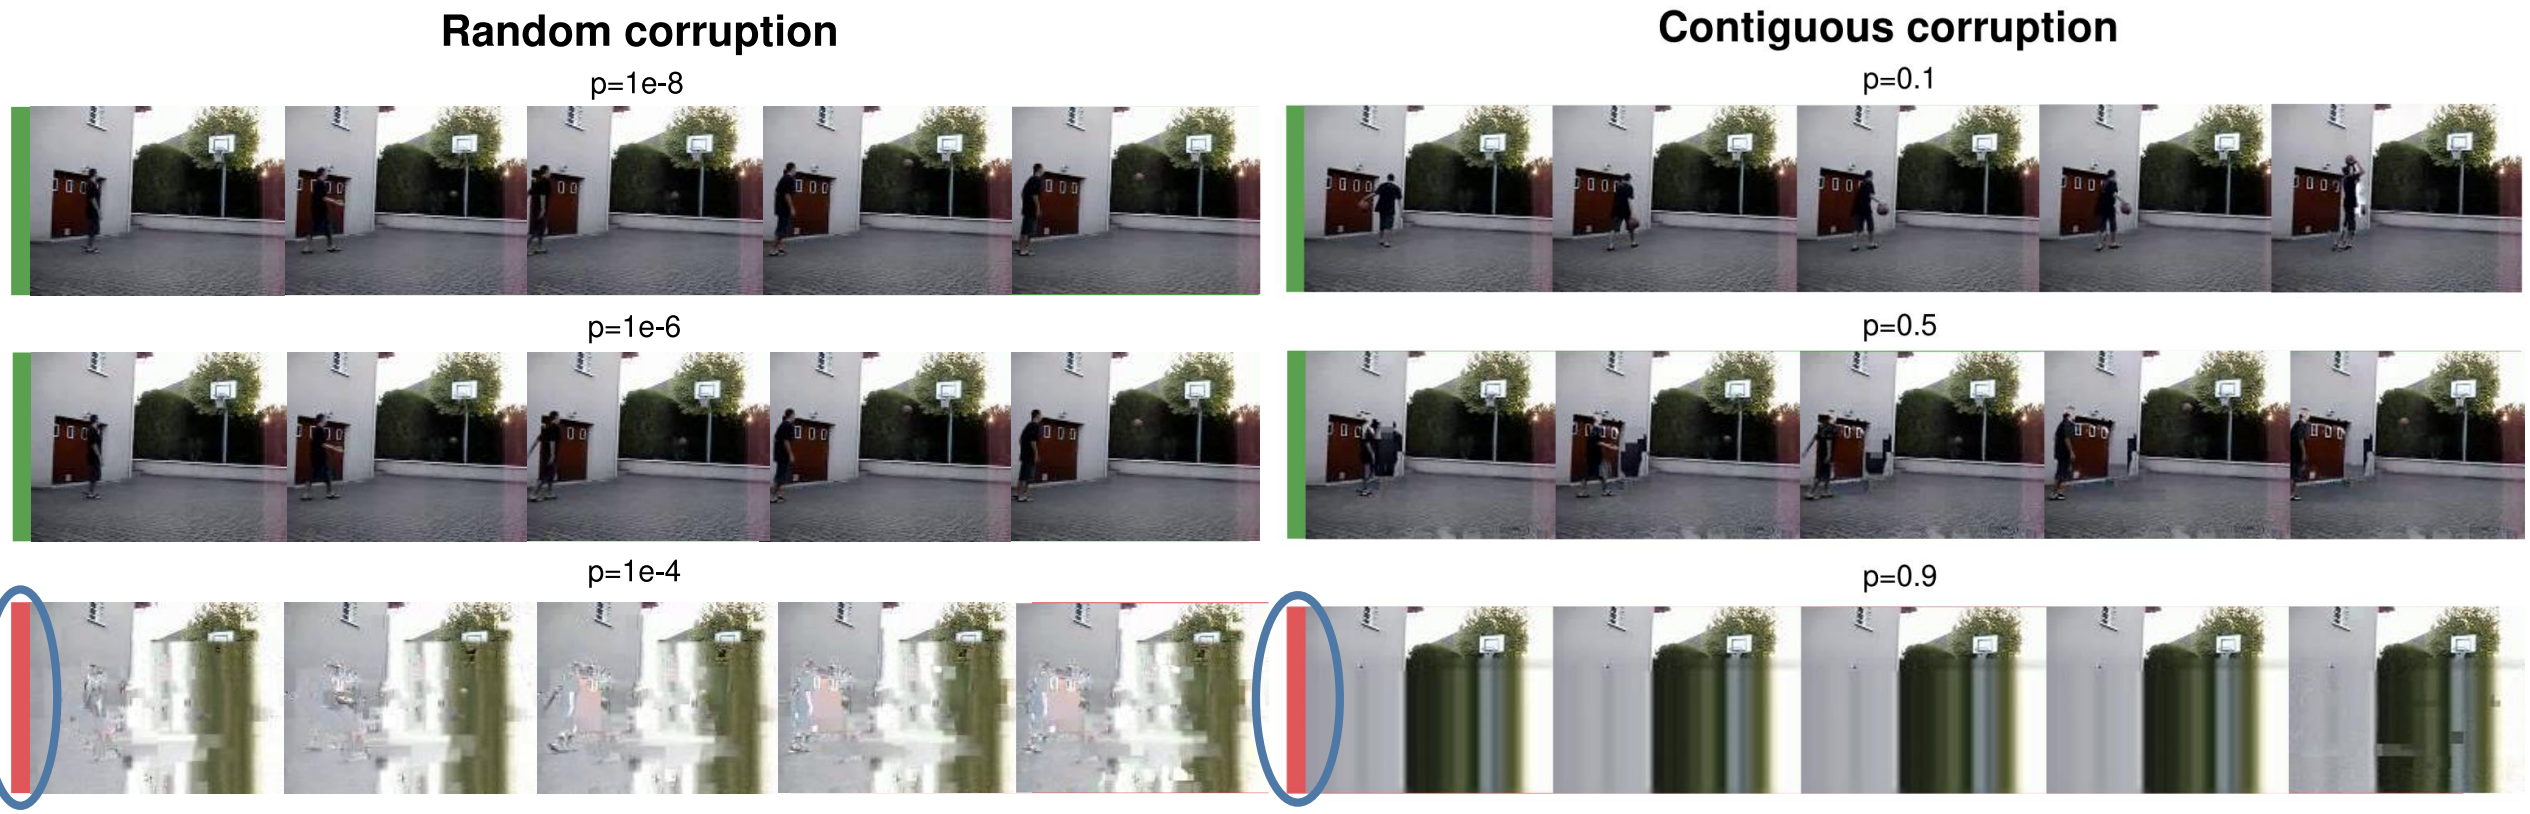

Clips that look worse (more corrupted) tend to be classified incorrectly

Model prediction: correct incorrect

Model errors correlate with more visually distorted clips given constant corruption proportion and strategy

Random corruption, p = 1e-4

ApplyEyeMakeup

✓ Correct

WalkingDog

X Incorrect

More visibly distorted = more likely to be incorrect

Model prediction: correct incorrect

Incorrectly classified examples are more quantitatively distorted under pixel-space Euclidean distance

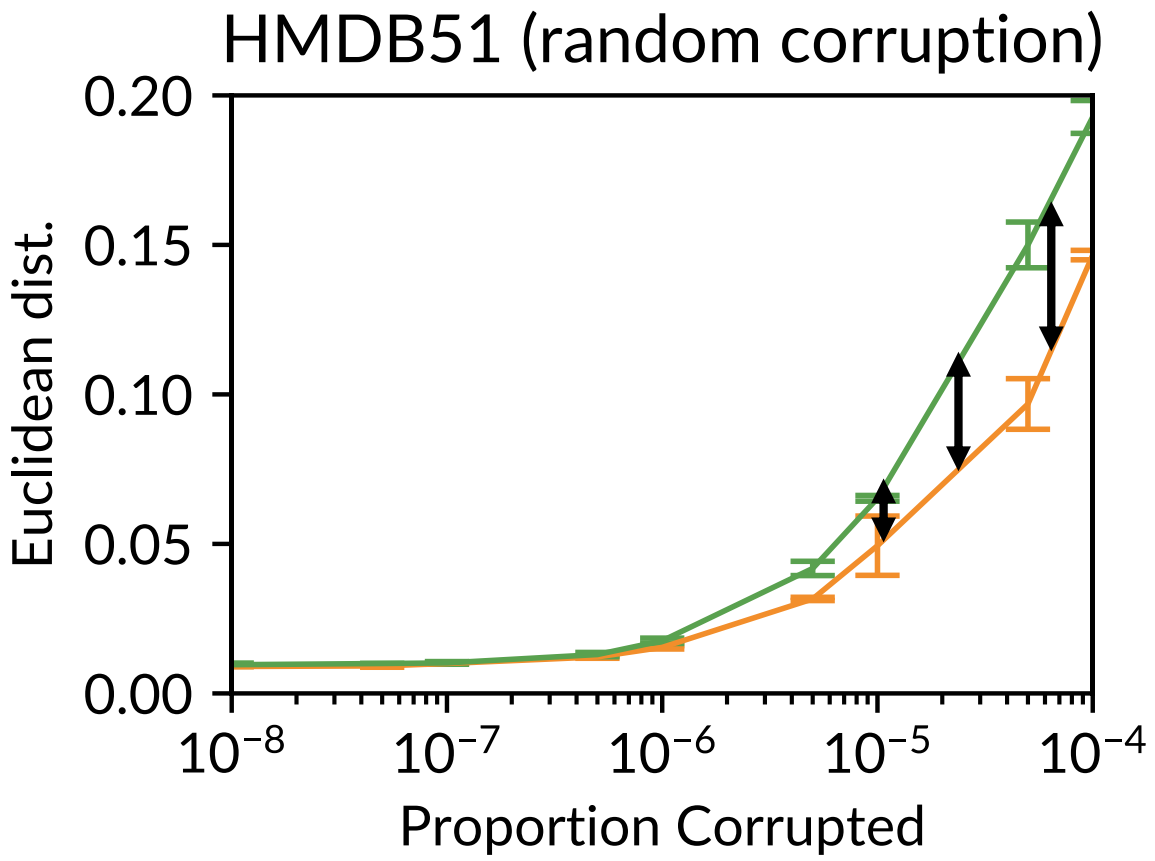

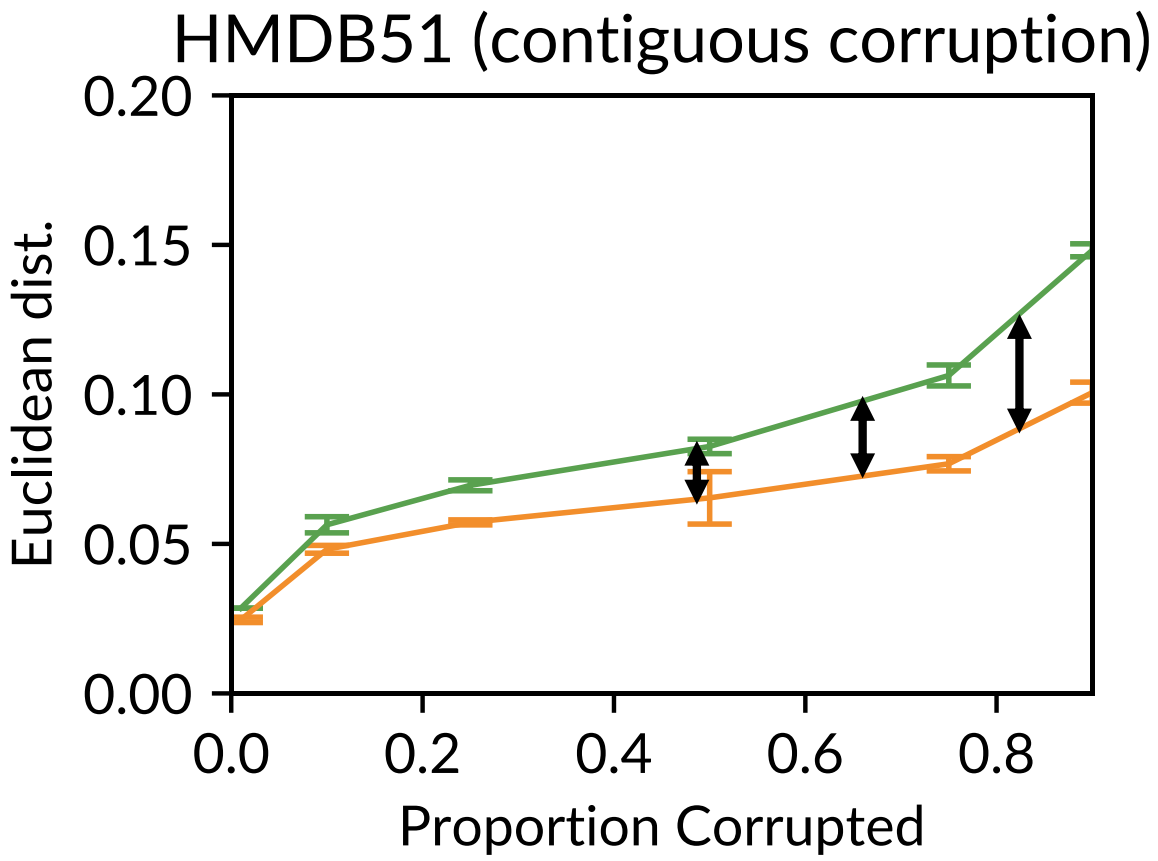

Summary

- As proportion of file corrupted goes up, accuracy goes down
- Clips that look more distorted tend to be classified incorrectly
- Clips that are more distorted under pixel-space Euclidean distance tend to be classified incorrectly

Contact: tchang97@cs.stanford.edu