Closing Remarks

Summary

- Top-5 accuracy has dropped
  - 1st place 79.25% (94.78% in 2017)
- Noise is the major concern
  - Resampling + Bootstrapping
  - Self-supervised pre-training
  - Confidence/Clustering based data cleaning
  - Multiple instance learning
- Is noise the only concern?
  - Class imbalance
  - Very hard classes
- Few cares of meta-data

Open Questions

- Definition of classes
  - Long tail distribution
  - What classes are representative?
- Number of classes
  - Are 5,000 calesses enough or too many?
- Larger scale
  - Shall we increase the dataset size more?
- More tasks
  - Webly supervised semantic segmentation / object detection
  - Video Recognition

Thank you!