Always Be Dreaming: A New Approach for Data-Free Class-Incremental Learning






James Smith[1] Yen-Chang Hsu[2] Jonathan Balloch[1] Yilin Shen[2] Hongxia Jin[2] Zsolt Kira[1]

[1]Georgia Tech, [2]Samsung Research America

2021 ICCV OCTOBER 11-17 VIRTUAL

RIPL Robotics Perception and Learning

SAMSUNG RESEARCH AMERICA

Georgia Tech Machine Learning

Introduction: Neural Networks Forget

- Typical learning with neural networks
  - Offline training
  - Shuffled inputs satisfy i.i.d. assumptions
- Continual Learning with neural networks
  - Tasks encountered sequentially
  - Catastrophic forgetting: overwrite information from past task(s)
- **Question**: how do we acquire new information without forgetting?

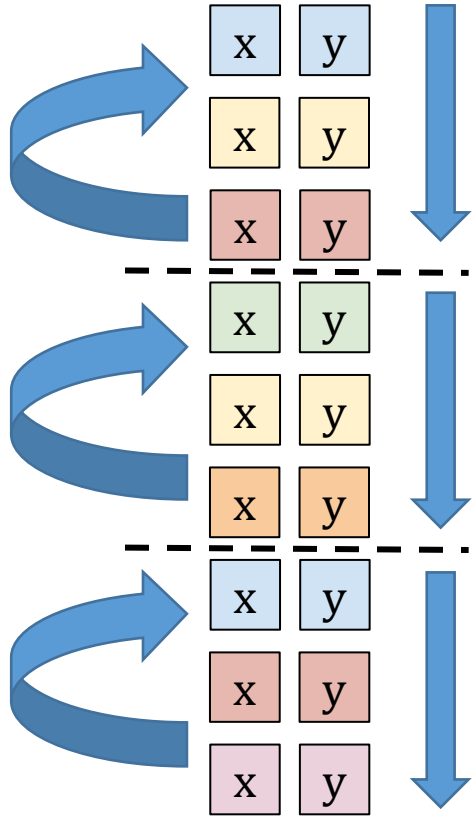

Replay in Continual Learning has Concerns

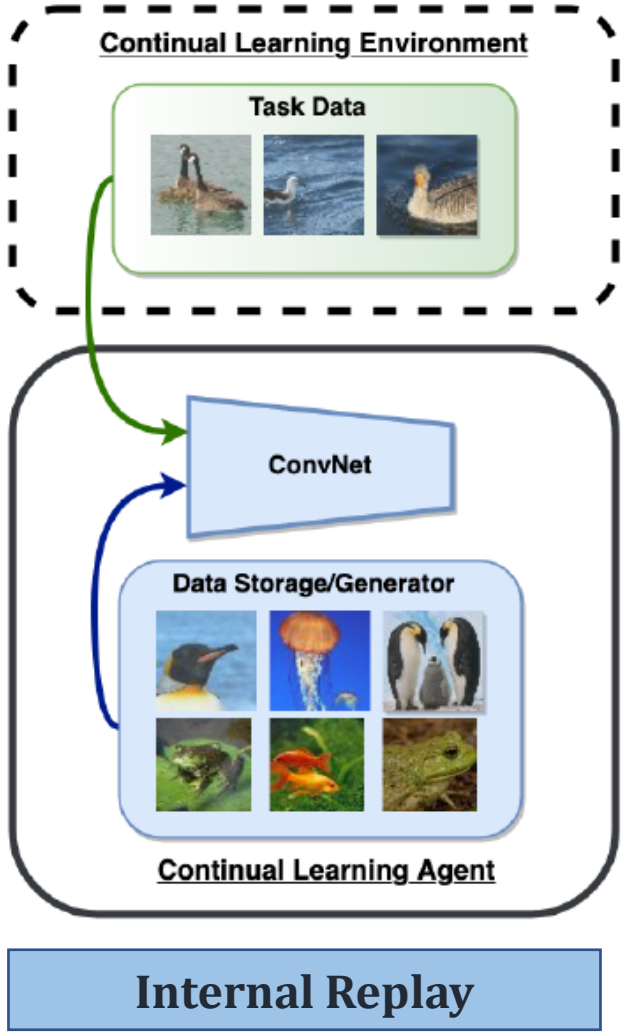

- **Internal Replay**
  - Store data from past tasks
  - Replay stored data in future tasks
- **Concerns**
  - Requires **substantial memory budget**
  - Some data cannot be stored due to **privacy concerns**

Idea: *Dream* Replay Images from Model

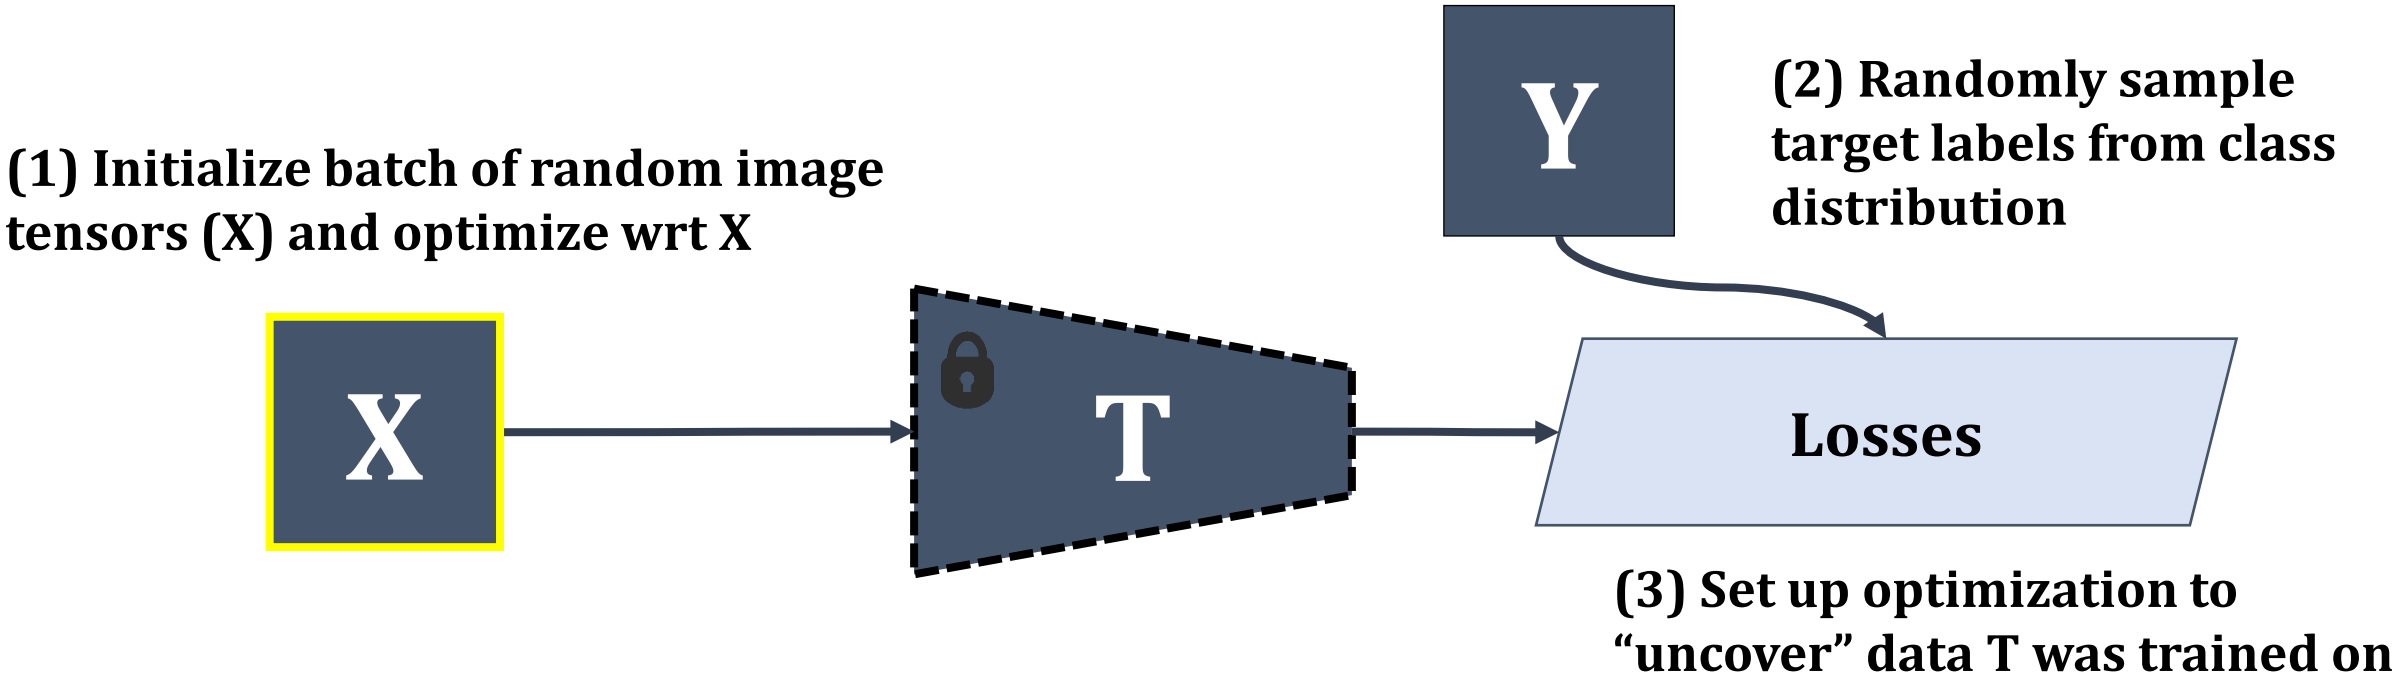

Problem: "dreamed" images fail for class-incremental continual learning distillation!

"**Dreaming to Distill: Data-Free Knowledge Transfer via DeepInversion**", Yin et al., CVPR 2020

Improve Distillation for *Dream* Images

- **Diagnosis**
  - Feature embeddings prioritize **domain** (i.e., synthetic versus real data) over **semantics** (i.e., task 1 versus task 2)
  - **Bias**: model predicts previous task images with current task labels!

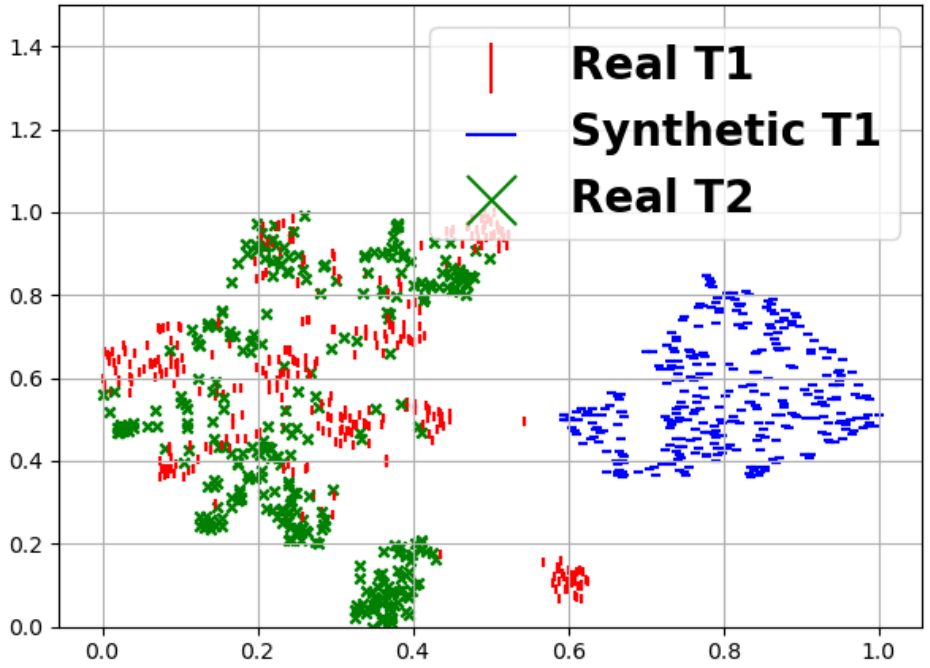

▲ Fig: t-SNE visualization for DeepInversion (Yin et al., 2020) on two iCIFAR-100 task

"**Dreaming to Distill: Data-Free Knowledge Transfer via DeepInversion**", Yin et al., CVPR 2020

Improve Distillation for *Dream* Images

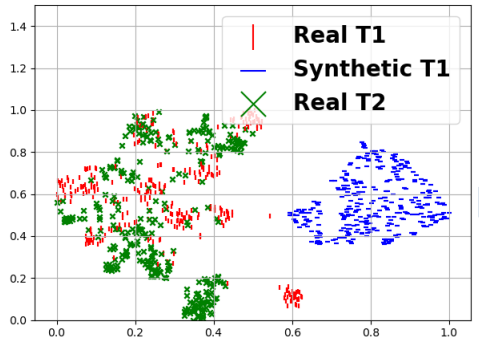

How can remove this bias problem?

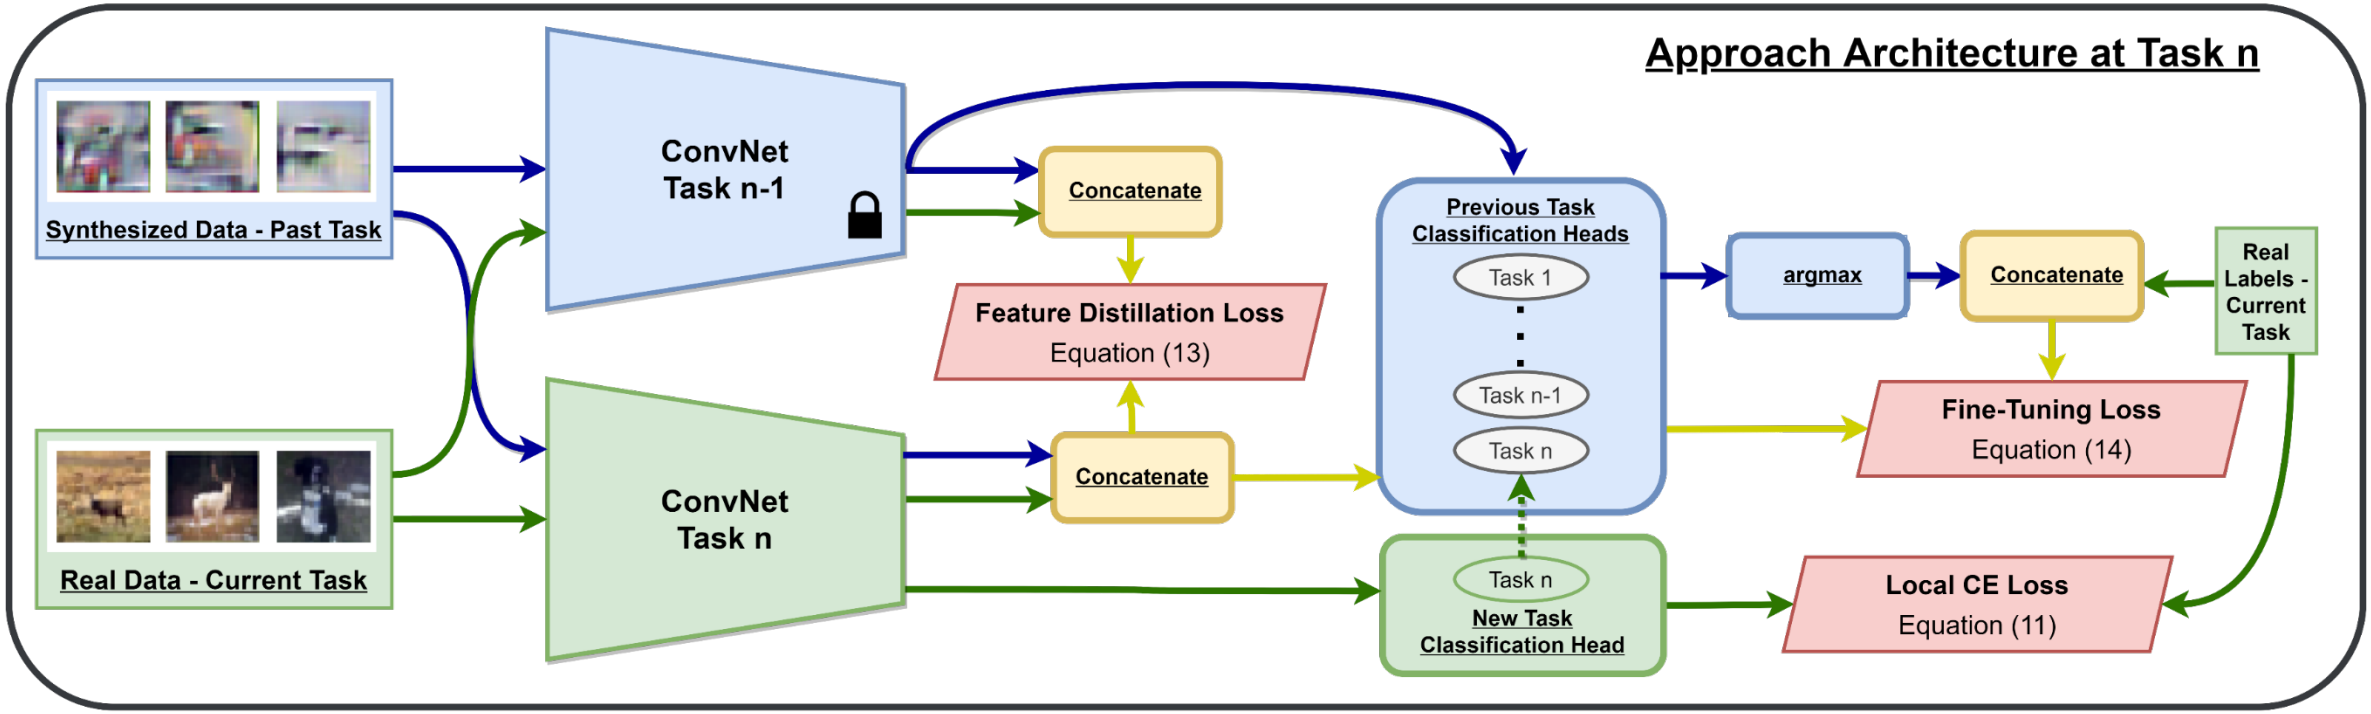

Improve Distillation for *Dream* Images

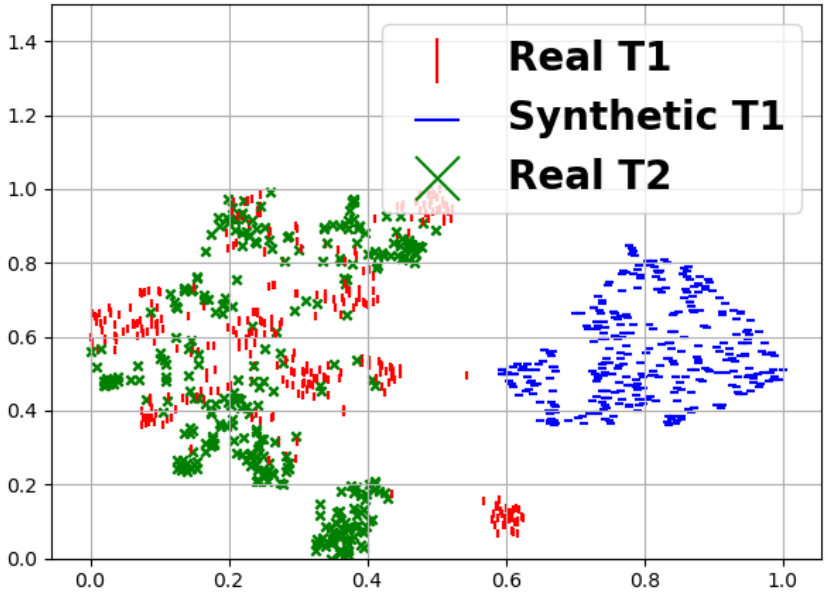

We remove this bias problem!

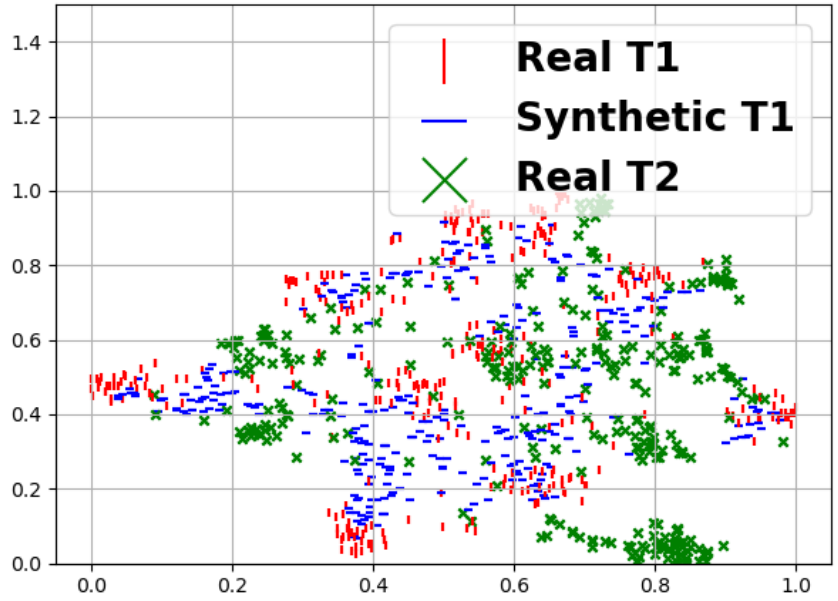

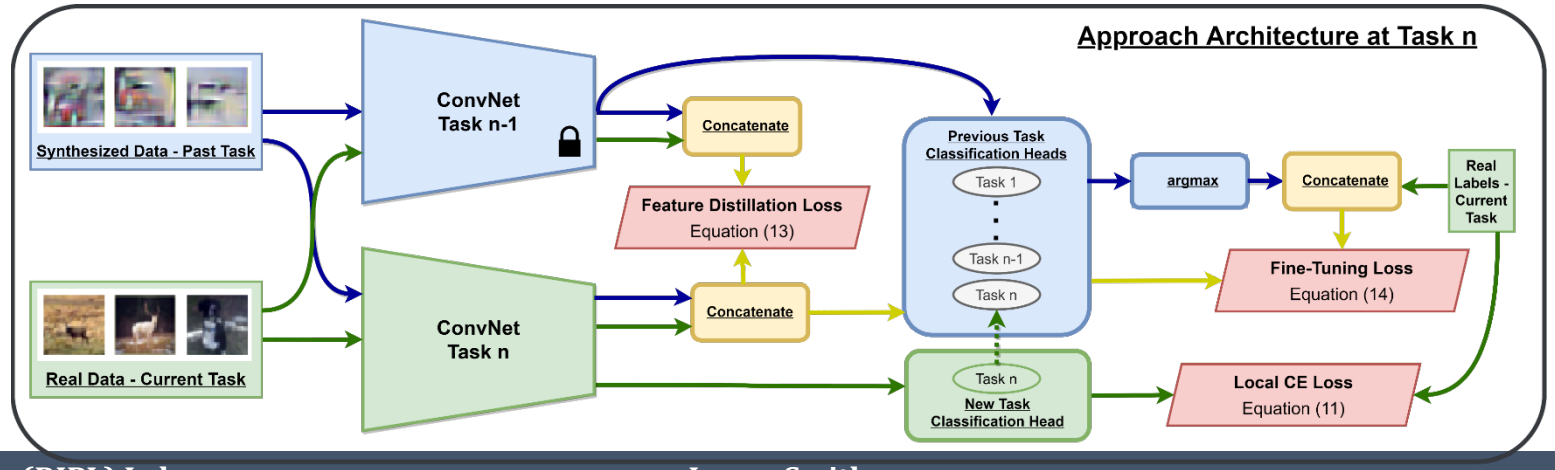

https://jamessealesmith.github.io/project/dfcil/






James Smith[1] Yen-Chang Hsu[2] Jonathan Balloch[1] Yilin Shen[2] Hongxia Jin[2] Zsolt Kira[1]

This work is supported by Samsung Research America.