Meta-Learning Representations for Continual Learning

Khurram Javed, Martha White
{kjaved,whitem}@ualberta.ca
amii · RLAI · University of Alberta

Overview

- Neural Networks trained online on a correlated stream of data suffer from catastrophic forgetting.
- We propose learning a representation that is robust to forgetting.
- To learn the representation, we propose **OML**, a second-order meta-learning objective that directly minimizes interference.
- Highly sparse representations naturally emerge by minimizing our proposed objective.

Motivation

- **Question:** Can we learn representations that are robust to catastrophic forgetting?

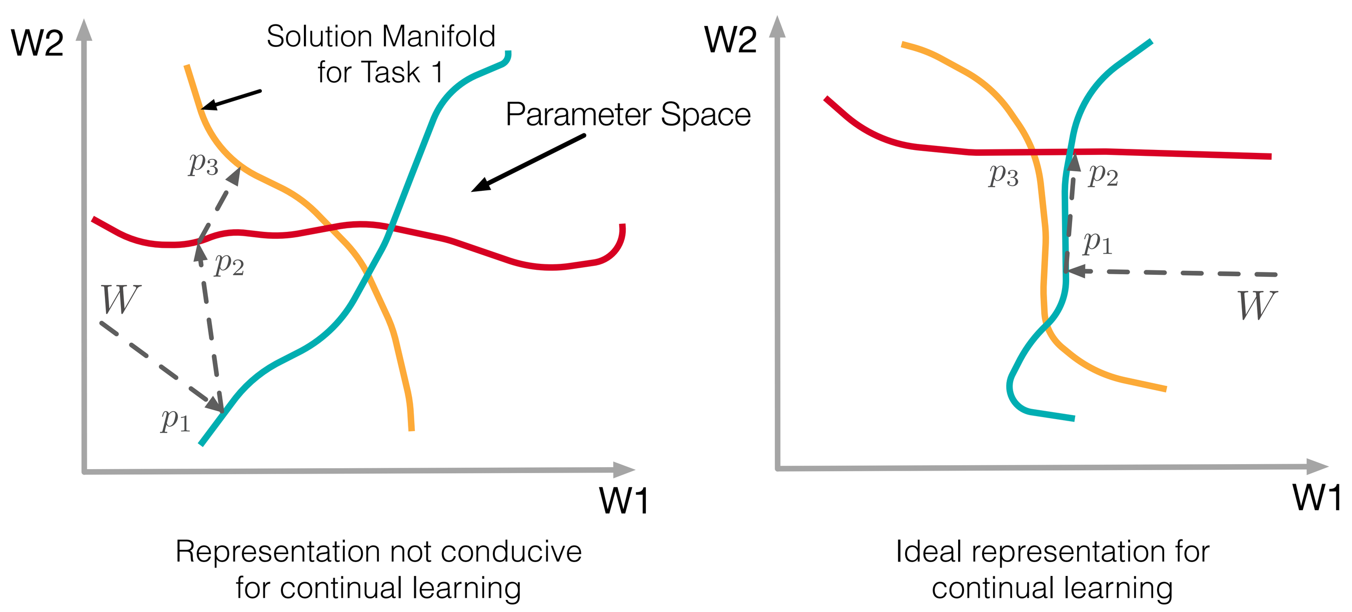

- **Hypothesis:** Some knowledge representations (right) are more conducive for continual learning than others (left).

Proposed architecture

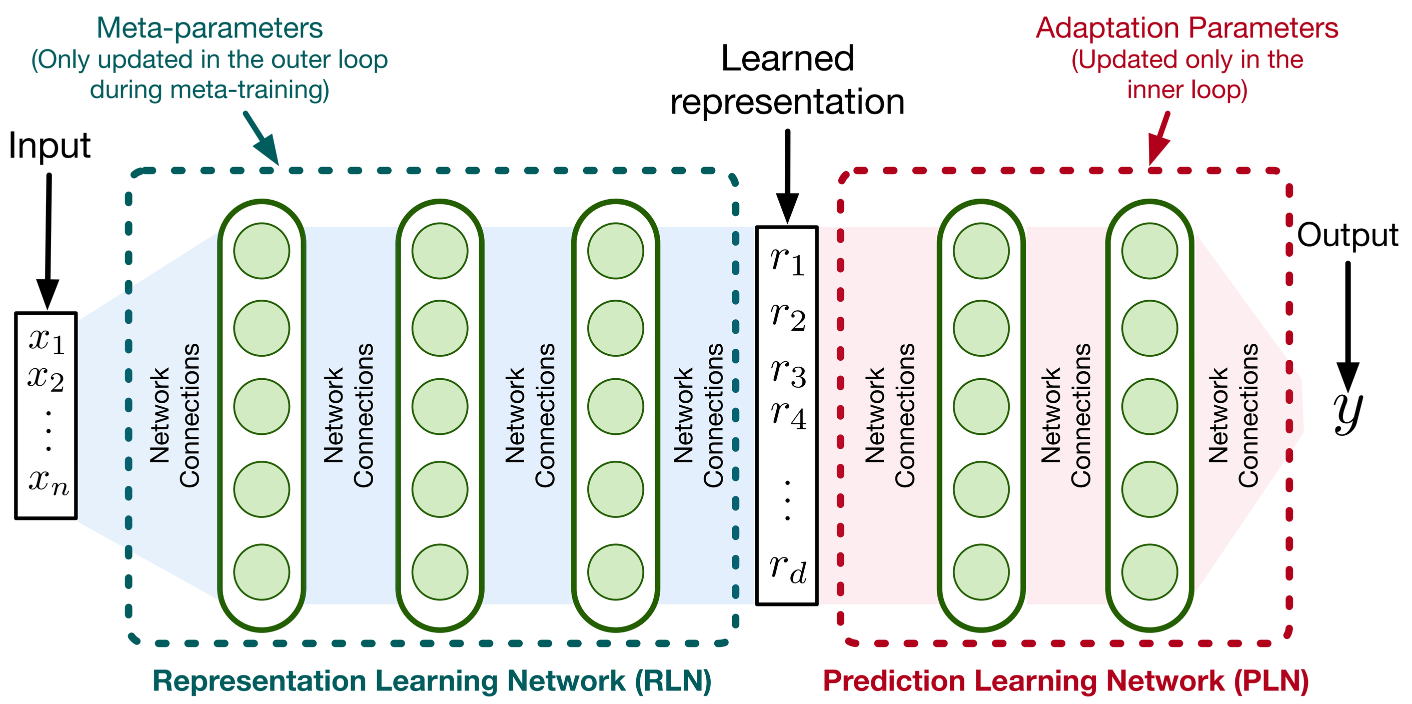

- **Meta-parameters:** A deep neural network that transforms high-dimensional input data to a representation $\mathcal{R}^d$ more conducive for continual learning.
- **Adaptation parameters** A simple neural network that learns continually from $\mathcal{R}^d$.

Meta-training

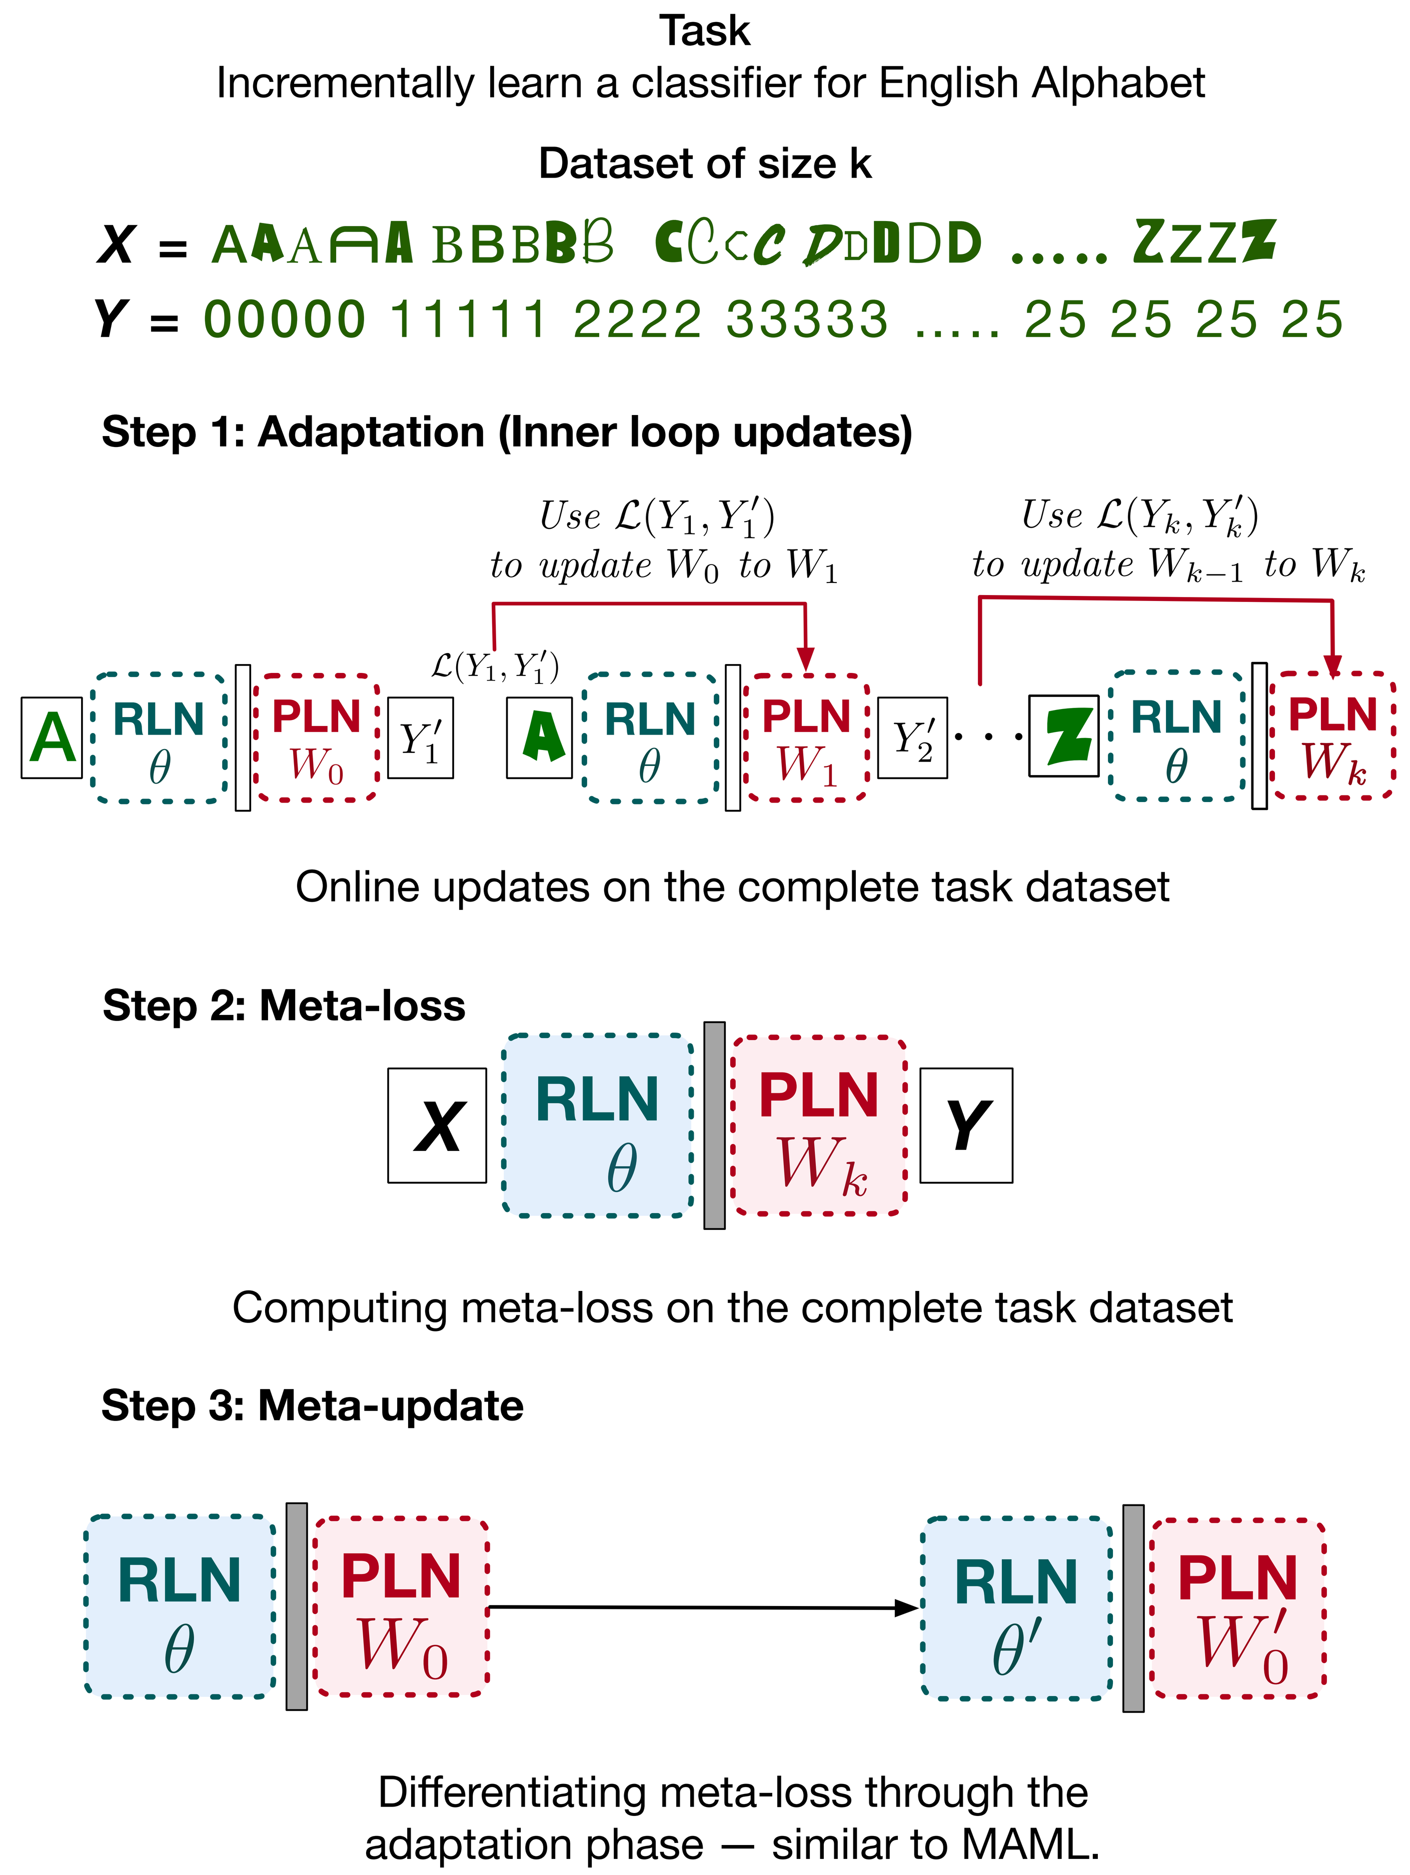

Meta-testing

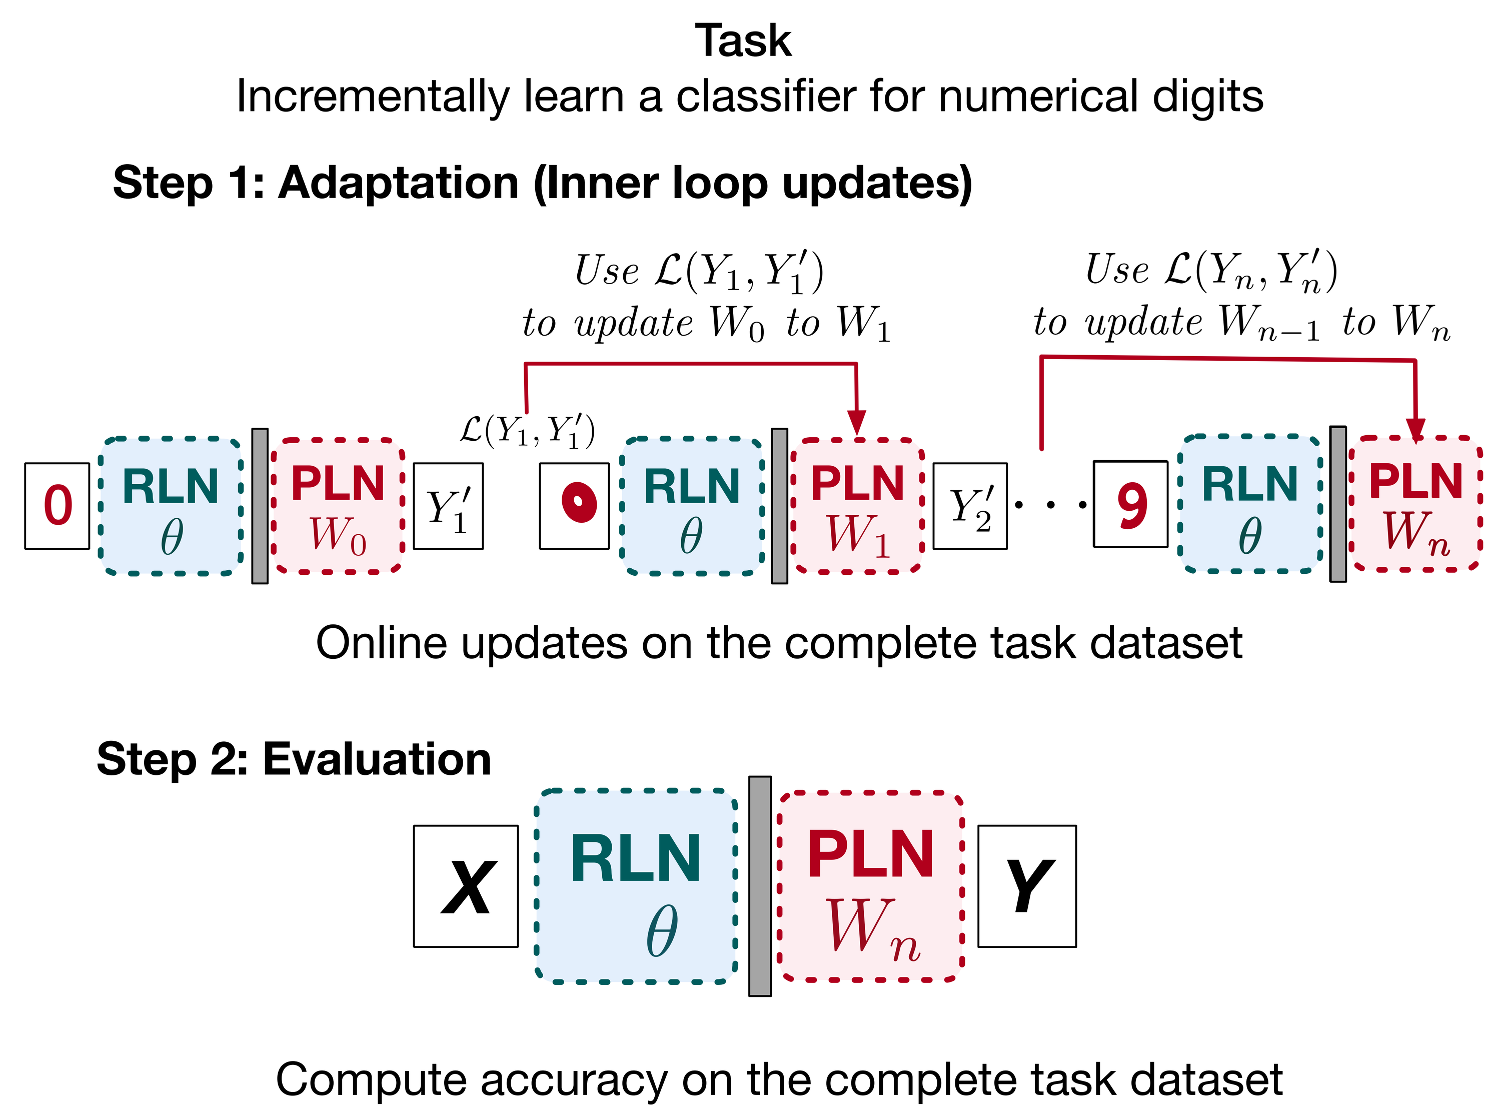

Results

We compare OML with a Pretraining, a method that learns a representation by pre-training on the meta-training dataset, MAML-Rep, a MAML like fast adaptation objective that also learns an RLN and SR-NN, a recent method that learns sparse representations.

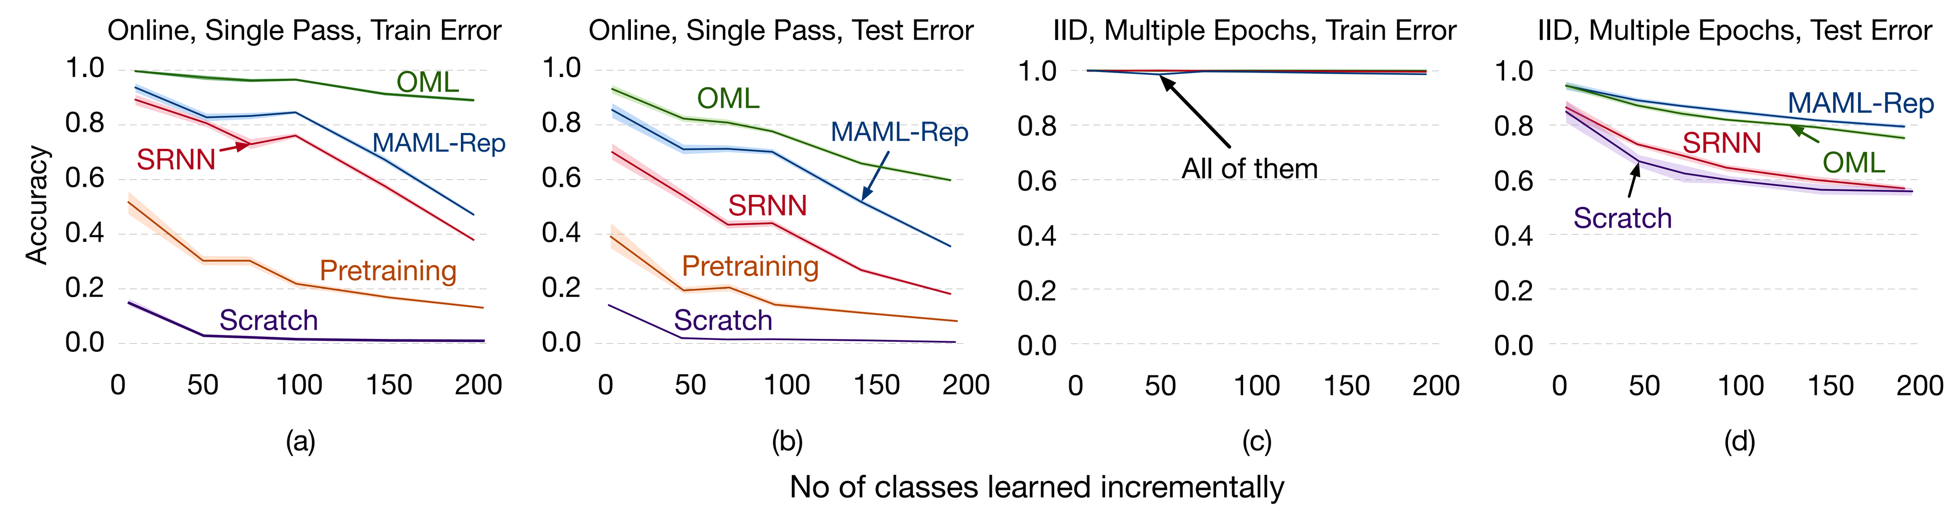

What sort of representations does OML learn?

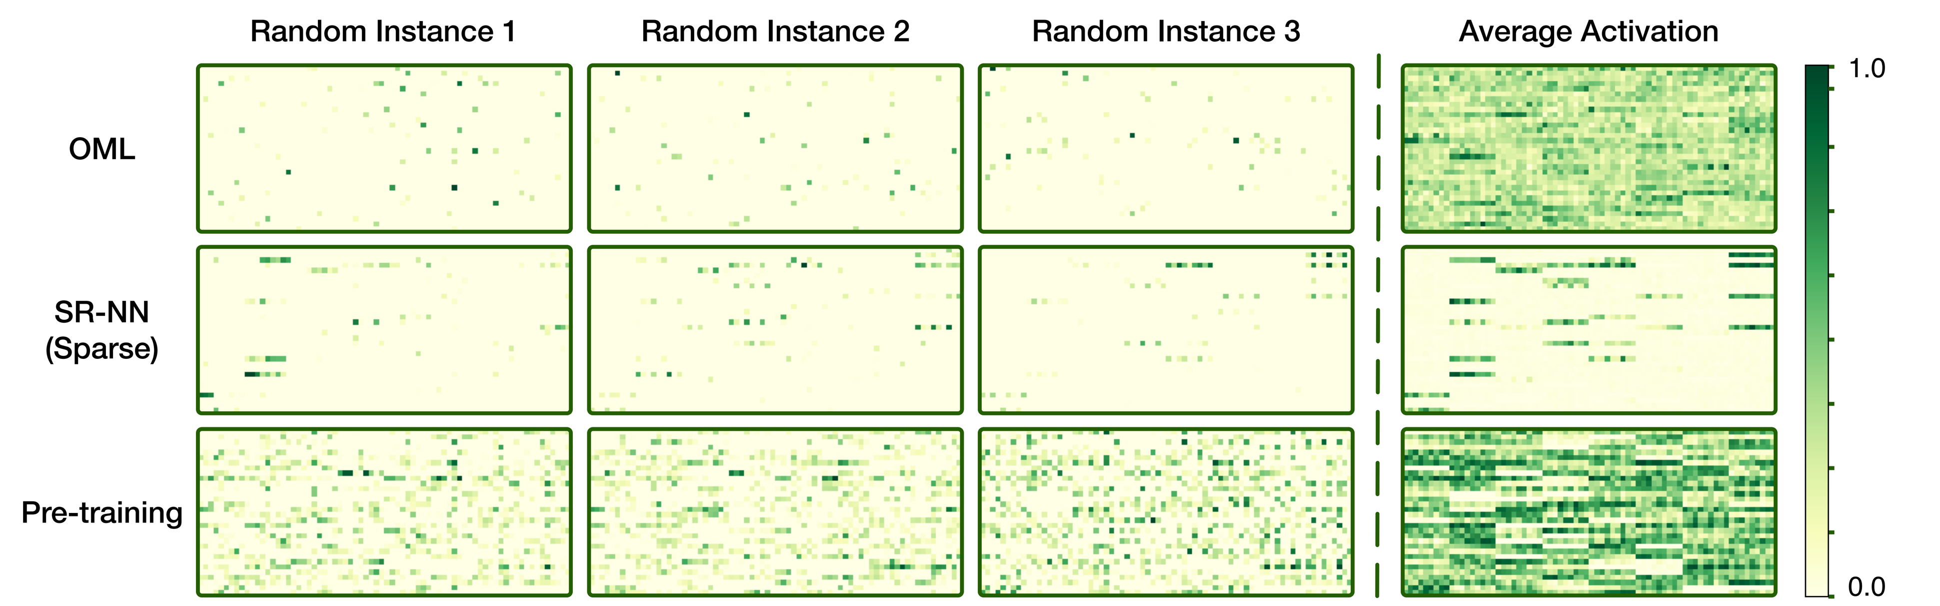

Figure: OML learns highly sparse representations without directly optimizing for sparsity. Moreover, unlike SR-NN, OML utilizes the complete representation space to represent different inputs.

Can we meta-learn a model initialization instead?

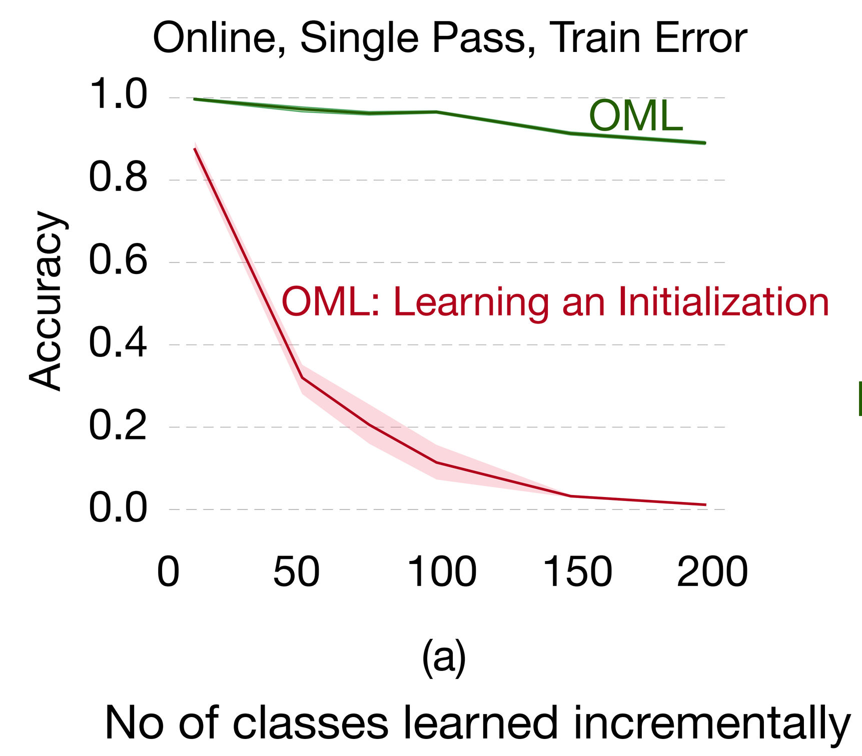

Answer: No!

Not effective when meta-testing involves hundreds of updates.

Future Work

- Continually meta-update the RLN as opposed to fixing it after meta-learning.
- Preliminary results show that it's possible to use OML to update representations online using an experience replay buffer. This can extend OML to more exciting settings, such as reinforcement learning.