



ARCADe: A Rapid Continual Anomaly Detector

Ahmed Frikha, Denis Krompaß & Volker Tresp

Siemens AG & Ludwig Maximilian University of Munich

01.2021 @ICPR2020

Introduction



- Continual Anomaly Detection (CAD) is unexplored
- One model for multiple sequentially learned AD tasks
- Challenges: Catastrophic forgetting & Overfitting to the majority class
- Learn a learning strategy that can cope with these challenges
- ARCADe is a first and strong baseline for CAD

A Robust Continual Anomaly Detector ARCADe – Meta-Training



SIEMENS
Ingenuity for life

Algorithm 1 ARCADe Meta-training Procedure

Require: D_{tr} : Set of meta-training task-sequences
Require: β : Learning rate for the meta-update
Require: K : Adaptation set size

- 1: Randomly initialize model parameters θ and parameter-specific learning rates α
- 2: **while** not done **do**
- 3: Sample a batch of task-sequences S_i from D_{tr}
- 4: Initialize meta-learning loss $L_{meta} = 0$
- 5: **for each** sampled S_i **do**
- 6: Adapt model initialization θ to S_i using the train set of each task and the learning rates α , yielding $\theta'_{i,J}$
- 7: Evaluate the adapted model $f_{\theta'_{i,J}}$ on S_i using the validation set of each task to compute $L_{s,i}$
- 8: $L_{meta} = L_{meta} + L_{s,i}$
- 9: **end for**
- 10: Update (θ, α) : $(\theta, \alpha) \leftarrow (\theta, \alpha) - \beta \nabla_{(\theta, \alpha)} L_{meta}$
- 11: **end while**
- 12: **return** Meta-learned parameters θ and learning rates α

Algorithm 2 ARCADe Meta-testing Procedure

Require: θ : Meta-learned initialization
Require: α : Meta-learned parameter-specific learning rates
Require: S_{test} : A meta-testing task-sequence

- 1: Adapt model initialization θ to S_{test} using the train set of each task and the learning rates α , yielding θ'_{J}
- 2: Evaluate the adapted model $f_{\theta'_{J}}$ on S_{test} using the validation set of each task included it

Results – Comparison to baselines



SIEMENS
Ingenuity for life

Fig. 1. Retained accuracy on Omniglot

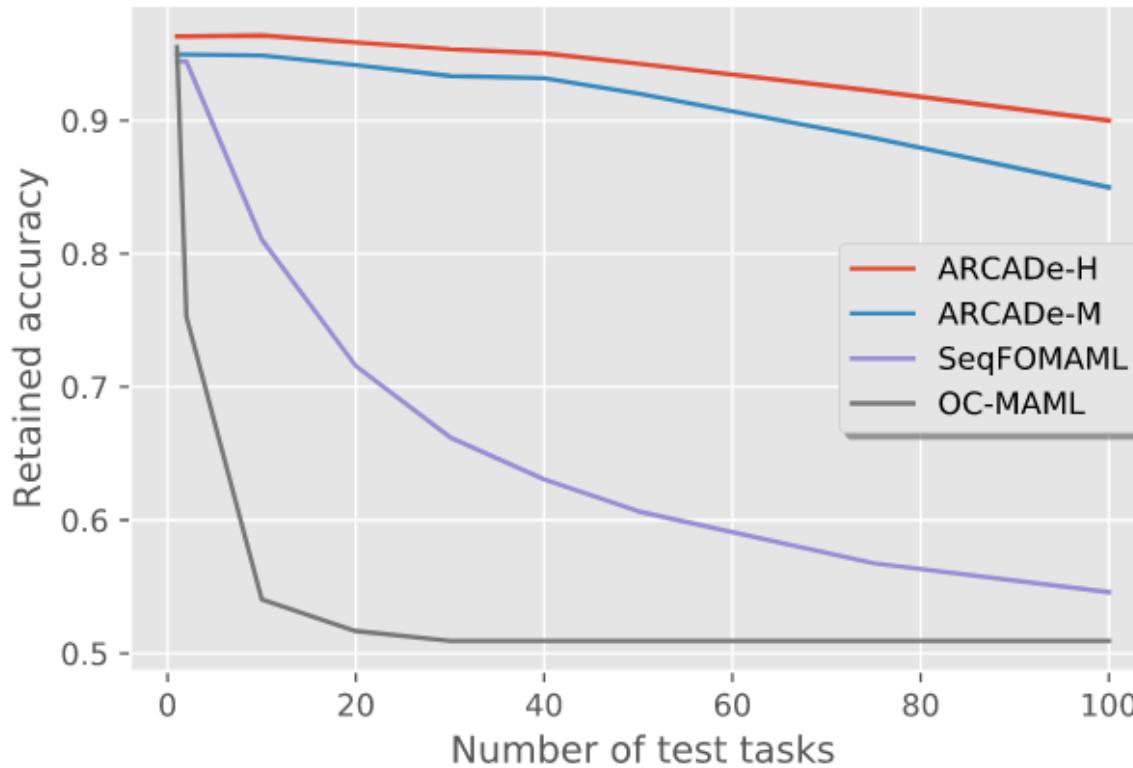
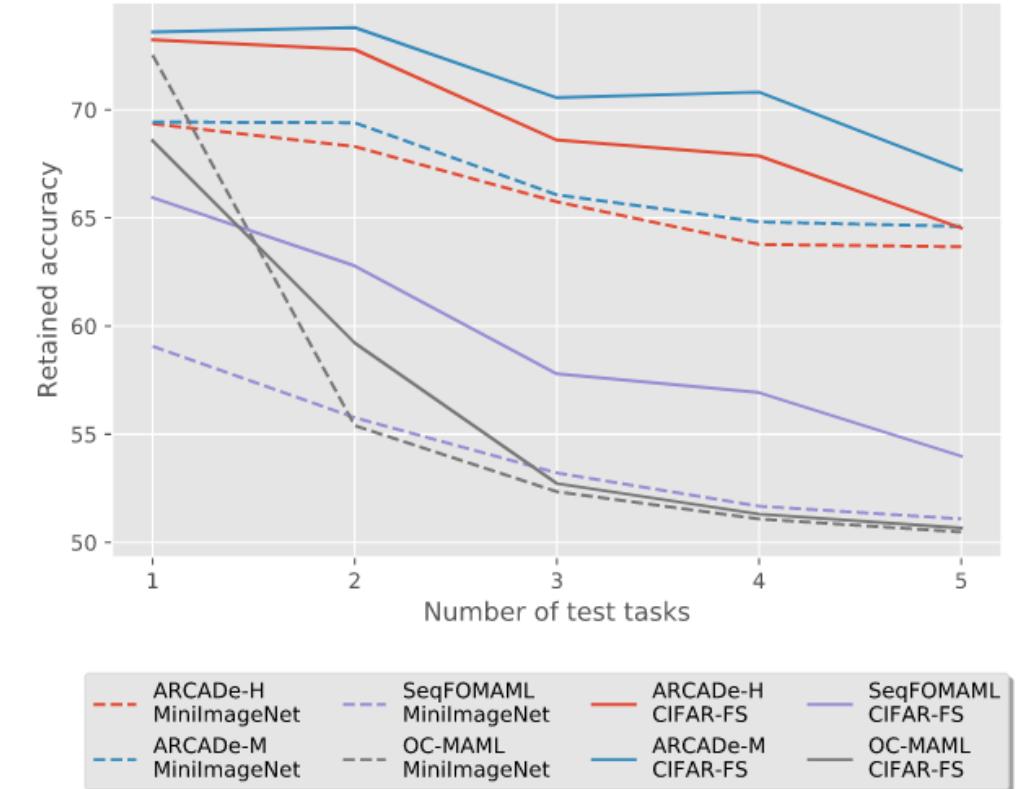


Fig. 2. Retained accuracy on MiniImageNet and CIFAR-FS



Results – The resulting learning strategy



SIEMENS
Ingenuity for life

Fig. 3. Layer-wise mean and percentage of positive learning rates meta-learned by ARCADe-M

