Cross-domain few-shot image classification
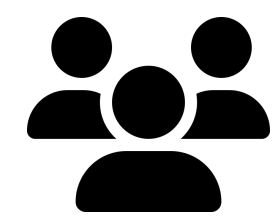
Pre-trained backbones
+ fine-tuning
+ batch learning
+ reduced domain gap

>

"de novo" training
+ episodic learning

Results of the NeurIPS'22 Cross-Domain MetaDL Competition

Dustin Carrión-Ojeda, Ihsan Ullah, Sergio Escalera, Isabelle Guyon, Felix Mohr, Manh Hung Nguyen, and Joaquin Vanschoren

1. INTRODUCTION

Traditionally, the few-shot learning evaluation protocols are within domain. But, our NeurIPS'22 Cross-Domain MetaDL Competition challenged participants to generalize across domains while comparing "de novo" training with the use of pre-trained backbones.
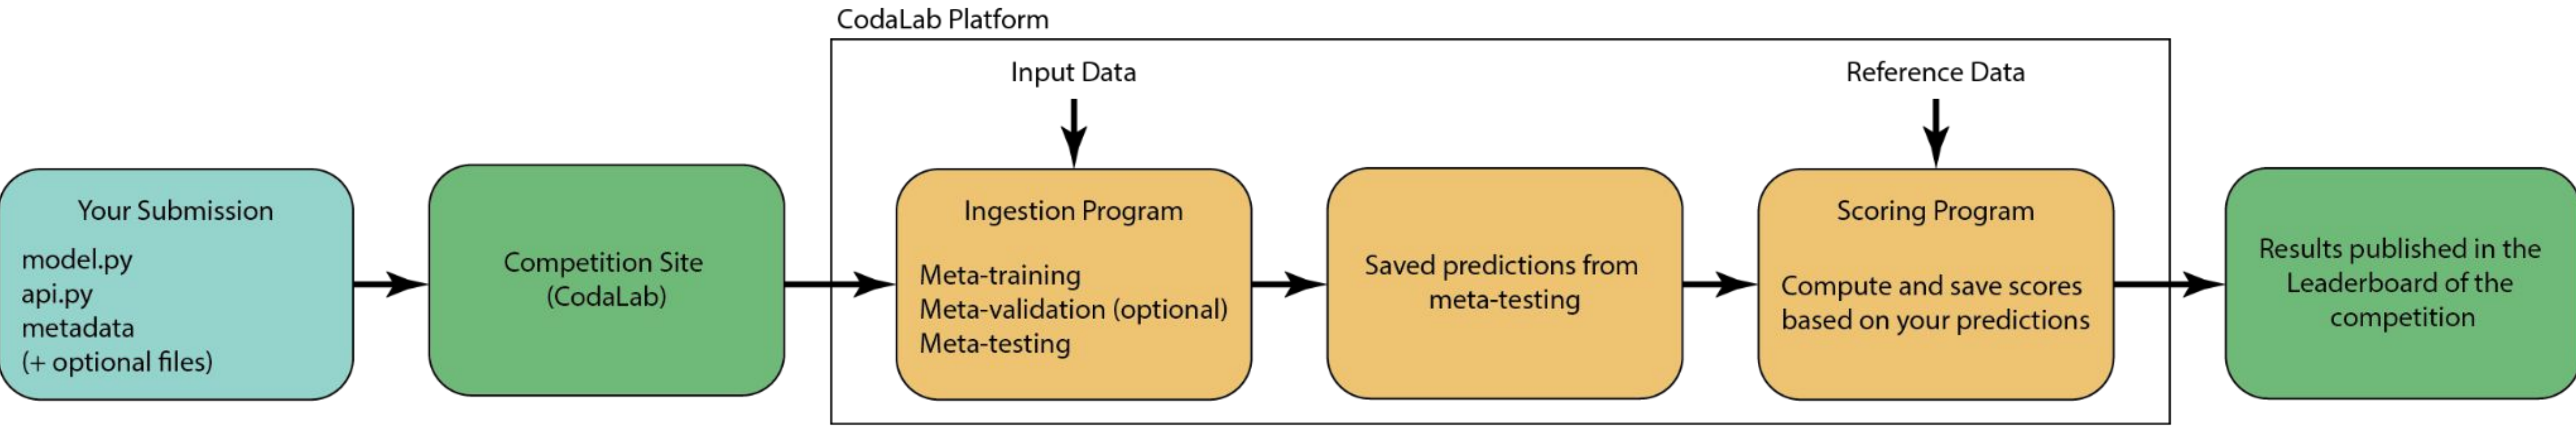
2. METHODS

- **Problem setting:** The meta-test tasks came from <u>10 domains</u> and were <u>any-way</u> ($\in [2, 20]$) <u>any-shot</u> ($\in [1, 20]$) with 20 unlabeled examples per class in the query set.
- **Data:** <u>Meta-Album</u> (30 datasets from 10 domains).
- **Baseline methods:** Train-from-scratch, Fine-tuning, Matching Networks, Prototypical Networks, FO-MAML, MetaDelta++.

3. WINNERS OVERVIEW

- **MetaBeyond:** Multiple lightweight task-adaptation modules.
- **EmmanuelPintelas:** Ensemble of Distance and Linear-based classifiers.

- **CDML:** Ensemble of pre-trained backbones + contrastive learning.
- **metaCD^2:** MetaDelta++ + contrastive learning + knowledge distillation.
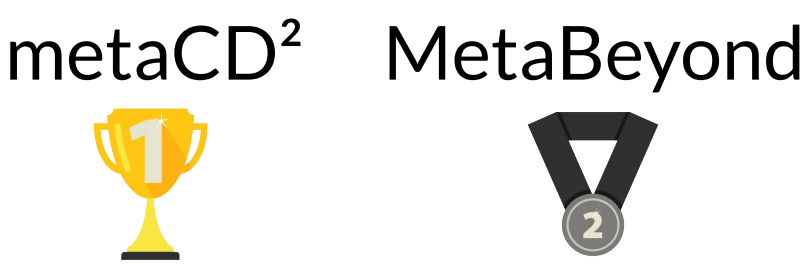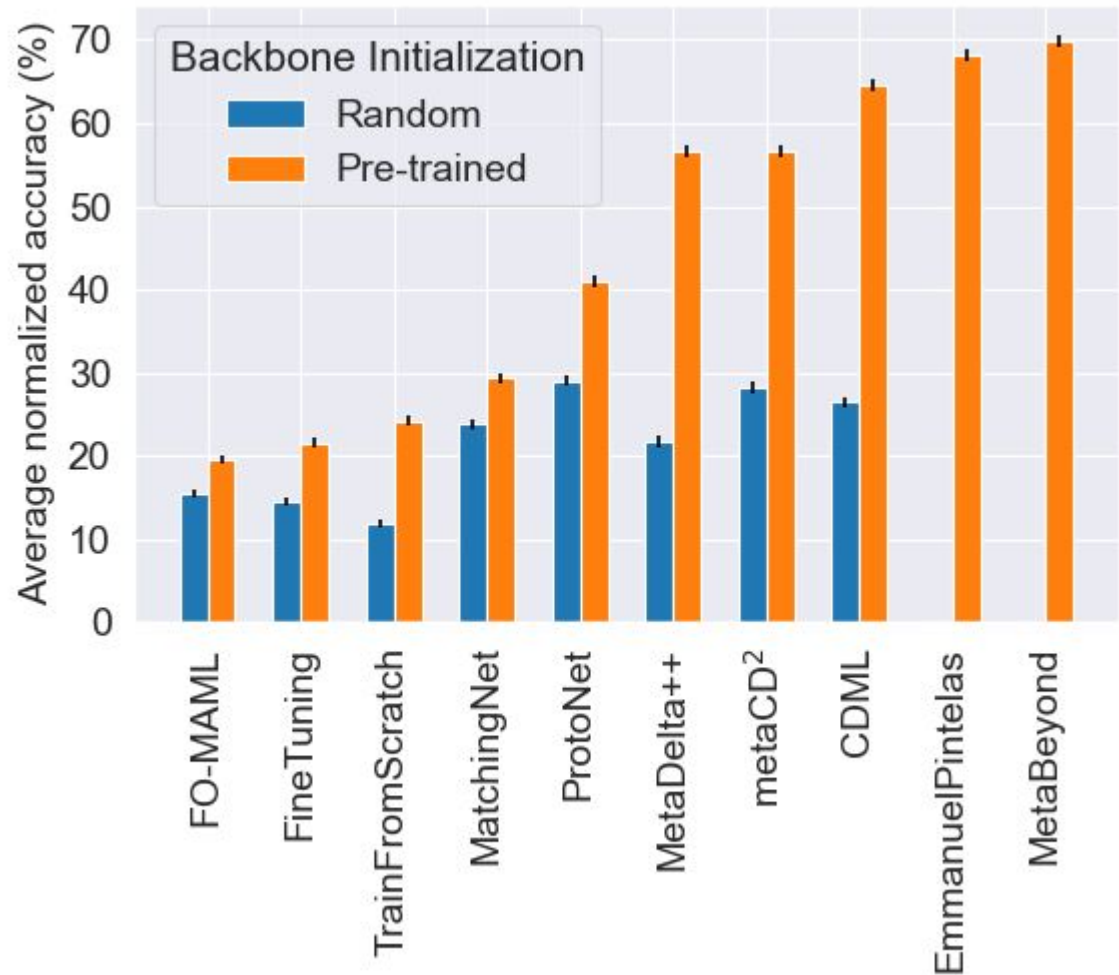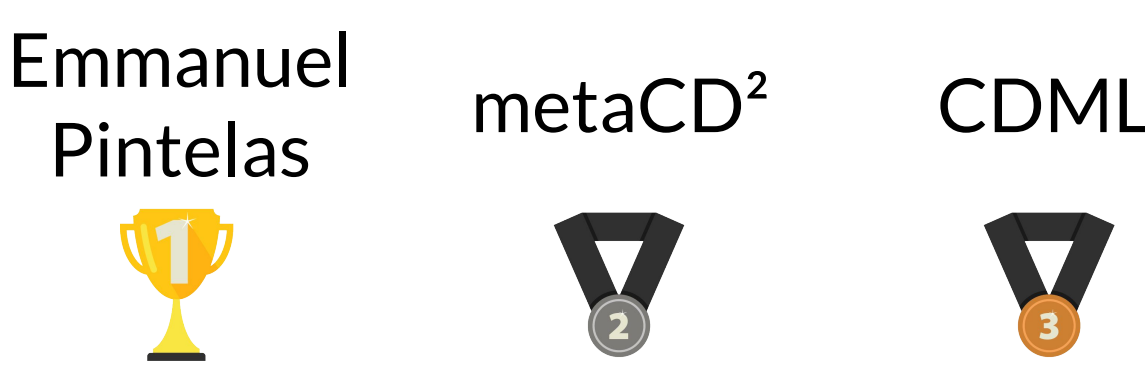
4. WINNERS PER LEAGUE

- **Free-style**

  MetaBeyond — 1
  Emmanuel Pintelas — 2
  CDML — 3

- **Meta-learning**

  metaCD^2 — 1
  CDML — 2

- **New-in-ML**

  metaCD^2 — 1
  MetaBeyond — 2

- **Women**

  MetaBeyond — 1

- **Rarely represented country**

  Emmanuel Pintelas — 1
  metaCD^2 — 2
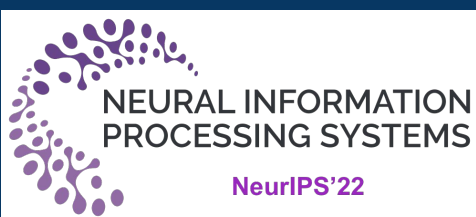 CDML — 3

Average normalized accuracy over 6 000 any-way any-shot meta-test tasks with 95% CIs computed at task level.
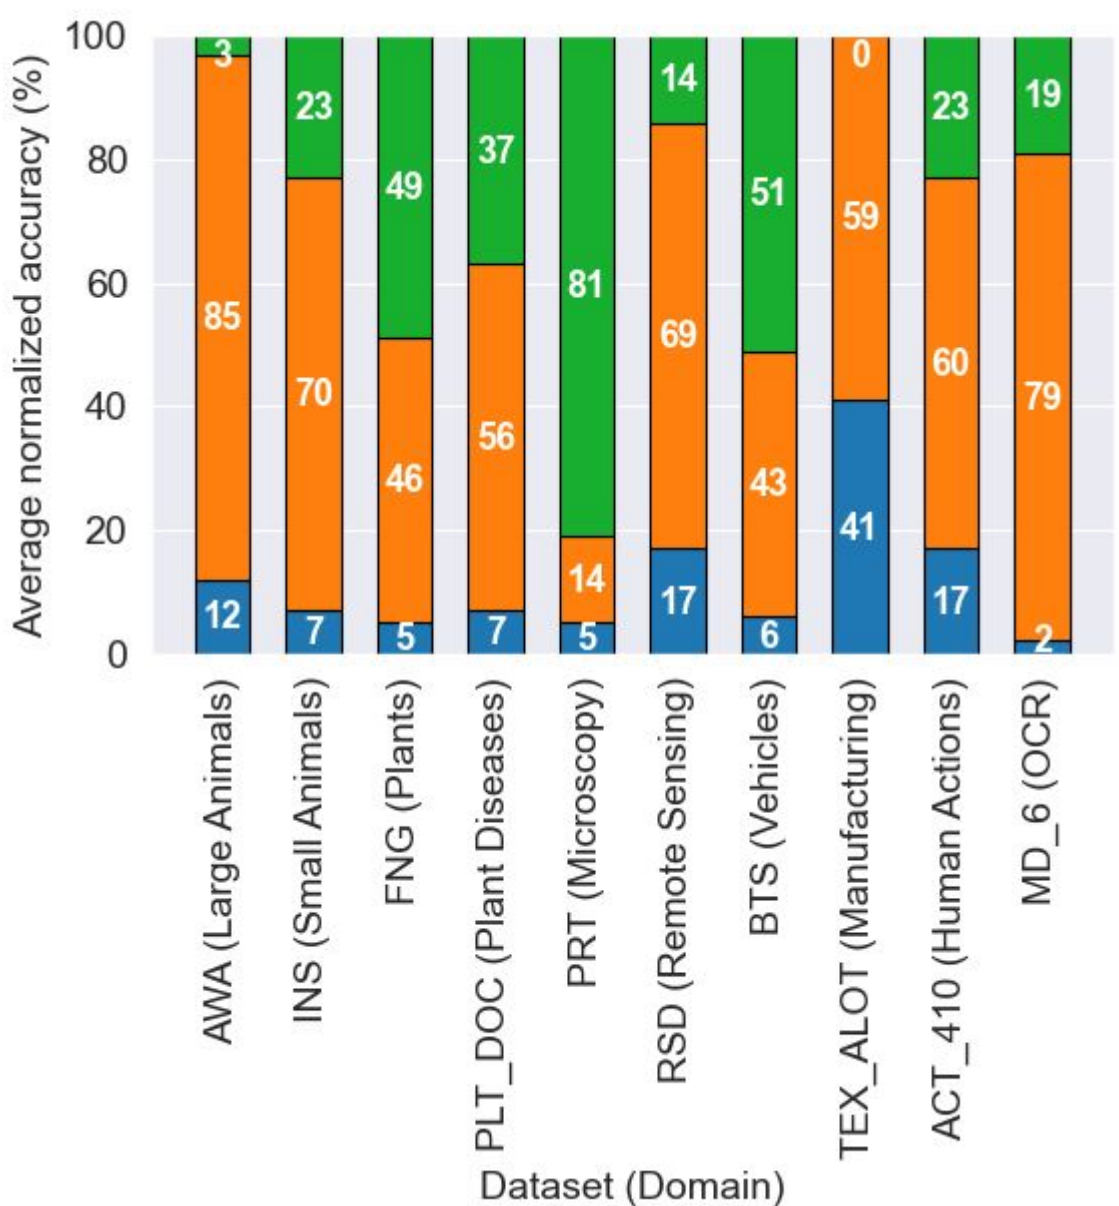
Average normalized accuracy computed over 600 any-way any-shot meta-test tasks per dataset. The top of the blue, orange, and green bars indicate the performance of TrainFromScratch w/o pre-training, MetaBeyond, and max., respectively.

5. COMPETITION EVALUATION FLOW