Mircea Cimpoi (he/him)

Google Scholar: 3SaQHMQAAAAJ |GitHub: mcimpoi | LinkedIn: mirceacimpoi

Summary

Research Engineer, with over 10 years of experience, specializing in infrastructure, data curation, and Computer Vision. Creator of the Describable Textures Dataset (DTD; CVPR2015), a widely used benchmark now standard datasets in Tensorflow and torchvision.
Experience with large scale codebases at Meta and Microsoft.

Employment

2025/10 - present – Self-paced study / career break
- deeplearning.ai courses / JAX / Rust, Kaggle AI Agents course, Advent of Code

2024/01 - 2025/09 – *AI Engineer*, **Sony AI**, Zurich, Switzerland
- Co-authored ***article*** in **Nature** on ethical collection of human-centric datasets:
  - Led the mitigation of fraudulent and suspicious patterns in the data collection process.
  - Reduced evaluation results storage by 100x (layered RLE) and codebase size by 50%.

2020/01 - 2023/07 – *Research Engineer*, **Meta Platforms**, Zurich, Switzerland
- Internal tooling for benchmarking CV and ML algorithms, for ad-hoc and periodic evaluations:
  - Reduced launch time from ~10 minutes to seconds; thousands of daily runs.
  - Exported and optimized models for on-device execution (quantization, TorchScript).

2017/09 - 2019/11 – *Postdoctoral Researcher*, **Czech Technical University (CIIRC)**, Prague, Czech Republic
- Large scale localization and navigation in *realistic* indoor spaces: InLOC, Matterport3D, Habitat - **NVIDIA GPU Seed Grant**. Co-authored: **1x CVPR, 1x NeurIPS, 2x TPAMI**.
- Collected ***Indoor Venues Dataset***, geographically diverse dataset of indoor scenes.

2016/02 - 2017/05 – *Machine Learning Engineer*, **Onfido (acquired in 2024, 400+ million)**, London, UK
- Automated the travel document classification and extraction pipeline (OCR, QR, MRZ).

2008/08 - 2011/07 – *Software Development Engineer*, **Microsoft**, Belgrade, Serbia
- *Shipped feature:* camera OCR, in Bing iPhone app, running offline, on iPhone 4GS.
- Designed storage format of individual character images for training OCR models (1M+ images).
- Microsoft Gold Star Award and 2x Patent awards for OCR and Document Layout Estimation.

Education

2011/10 - 2015/07 – Doctor of Philosophy, Engineering, **University of Oxford**
- Created Describable Textures Dataset (DTD) Tensorflow Datasets | Torch Vision Datasets.
- *Recognizing Describable Attributes and Materials of Textures in the Wild and Clutter* Project page.

Publications (Google Scholar 3SaQHMQAAAAJ)

[*NeurIPS 2018*][*TPAMI 2020*] I. Rocco, M. Cimpoi, R. Arandjelovic, A. Torii, T. Pajdla, J. Sivic, *Neighbourhood Consensus Networks* - **Spotlight (16%)**.
[*CVPR 2018*][*TPAMI 2019*] H. Taira, M. Okutomi, T. Sattler, M. Cimpoi, M. Pollefeys, J. Sivic, T. Pajdla, A. Torii - *InLoc: Indoor Visual Localization with Dense Matching and View Synthesis*, **Spotlight (6.7%)**.
[*CVPR 2015*][*IJCV 2016*] M. Cimpoi, S. Maji, A. Vedaldi, *Deep Convolutional Filter Banks for Texture Recognition and Segmentation*, **Oral (3.3%)**.
[*CVPR 2014*] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, A. Vedaldi, *Describing textures in the wild.*

Academic Service

Reviewer: CVPR: 2024; ICLR: 2021, NeurIPS: 2019 - 2022, 2024 (**top reviewer**), TPAMI: 2018