

Human Activity Recognition in an Open World (Abstract Reprint)

Derek S. Prijatelj¹, Samuel Grieggs², Jin Huang¹, Dawei Du³, Ameya Shringi³, Christopher Funk³, Adam Kaufman⁴, Eric Robertson⁴ and Walter J. Scheirer¹

¹Department of Computer Science and Engineering, University of Notre Dame

²Department of Mathematical and Computer Sciences, Indiana University of Pennsylvania

³Kitware, Inc., 1712 Route 9, Suite 300, Clifton Park, NY 12065, USA

⁴PAR Government, 421 Ridge St, Rome, NY 13440, USA

⁵Department of Computer Science and Engineering, University of Notre Dame

dprijate@nd.edu, sgrieggs@iup.edu, jhuang24@nd.edu, cvdaviddo@gmail.com, ashringi@google.com,
christopher.funk@kitware.com, awk11@outlook.com, robertsone361@gmail.com,
walter.scheirer@nd.edu

Abstract Reprint. This is an abstract reprint of a journal article by [Prijatelj *et al.*, 2025].

Abstract

Managing novelty in perception-based human activity recognition (HAR) is critical in realistic settings to improve task performance over time and ensure solution generalization outside of prior seen samples. Novelty manifests in HAR as unseen samples, activities, objects, environments, and sensor changes, among other ways. Novelty may be task-relevant, such as a new class or new features, or task-irrelevant resulting in nuisance novelty, such as never before seen noise, blur, or distorted video recordings. To perform HAR optimally, algorithmic solutions must be tolerant to nuisance novelty, and learn over time in the face of novelty. This paper 1) formalizes the definition of novelty in HAR building upon the prior definition of novelty in classification tasks, 2) proposes an incremental open world learning (OWL) protocol and applies it to the Kinetics datasets to generate a new benchmark KOWL-718, 3) analyzes the performance of current state-of-the-art HAR models when novelty is introduced over time, 4) provides a containerized and packaged pipeline for reproducing the OWL protocol and for modifying for any future updates to Kinetics. The experimental analysis includes an ablation study of how the different models perform under various conditions as annotated by Kinetics-AVA. The code may be used to analyze different annotations and subsets of the Kinetics datasets in an incremental open world fashion, as well as be extended as further updates to Kinetics are released.

References

[Prijatelj *et al.*, 2025] Derek S. Prijatelj, Samuel Grieggs, Jin Huang, Dawei Du, Ameya Shringi, Christopher Funk, Adam Kaufman, Eric Robertson, and Walter J. Scheirer. Human activity recognition in an open world. *J. Artif. Int. Res.*, 81, January 2025.