

Learning from Visual Data in the Wild

Oisin Mac Aodha

University of Edinburgh

The wealth and complexity of visual information potentially observable by artificial systems deployed in the real world vastly exceeds the comparative simplicity of our carefully curated benchmark datasets. To operate safely and reliably in challenging environments, artificial systems need to be able to correctly recognize previously learned concepts, not confuse these known concepts with novel ones, and be able to differentiate novel concepts so that they can be grouped and efficiently learned. In this talk, I will discuss recent work from my group on problems related to learning from visual data in the context of biodiversity monitoring. Specifically, I will cover our recent work on self-supervised learning of shape from images, incremental category discovery, and how to learn species range maps to enhance perception.