

Tutorial Proposal for FLINS-ISKE2026 Conference

Title: Detecting Distribution Shifts in the Wild: Techniques and Future Perspectives

Duration: 1.5 Hours

Organizers: Bo Peng

Presenter: Bo Peng

¹ *University of Technology Sydney, Australia*

Email: Bo.Peng-7@student.uts.edu.au;

1 Introduction

Despite the significant progress in machine learning that has facilitated a broad spectrum of real-world tasks, existing models often operate under a closed-world scenario, where test data stems from the same distribution as the training data. Aiming at identifying unexpected inputs from unknown classes, out-of-distribution (OOD) detection has emerged as a pivotal approach to enhancing the reliability of machine learning models. This tutorial presents a comprehensive and in-depth survey of recent advances in OOD detection. We begin by introducing the tutorial goals and real-world motivations. Afterwards, we first focus on traditional OOD detection and subsequently delve into OOD detection with pre-trained vision-language representations. We conclude with a summary of current progress and a discussion of key challenges, open research questions, and future directions.

In summary, the objective and motivation of this tutorial is three-fold:

- present milestones in OOD detection
- review a broad selection of previous works in different categories
- uncover future directions of OOD detection.

2 Tutorial Outline and Schedule

The outline and schedule of this tutorial is given as follows:

1. **Opening, Background, and Key Insights**
 - Tutorial goals and real-world motivations.
 - Problem definition of OOD detection
2. **Traditional OOD Detection**
 - Post-hoc OOD detection
 - Outlier exposure for OOD detection
3. **zero-shot OOD Detection with Vision-language Models**
 - Introduction to vision-language models
 - Zero-shot OOD detection with pre-trained vision-language representations
4. **Conclusion and Future Trends**
5. **Q&A**

3 Teaching Methods

We promote the tutorial as follows:

- We plan to accompany our tutorial with intuitive examples.
- We will provide the participants with specific time slots to ask their questions.
- We will release all related materials (e.g., presentation slides, references, open-source data and code, etc.) in advance.

4 Target Audience and Prerequisites

OOD detection is a foundational topic in trustworthy machine learning. As such, our potential audience includes academic researchers and industry practitioners working in machine learning, data mining, computer vision, and natural language processing. Academic researchers may draw on established concepts and theoretical frameworks to develop novel methodologies and address emerging challenges. In contrast, industry practitioners are primarily concerned with translating theoretically grounded algorithms into reliable deployment of machine learning models for safety-critical applications. We assume that the audience has a general background in probability and statistics, along with a preliminary knowledge of deep neural networks.