



# Preliminary Program 2023

Speaker	Topic	Abstract	Date
<b>Week 1</b>			
<b>Shawe-Taylor, John</b>	The Foundations of Kernel Methods for Machine Learning		9 & 10 January 2023
<b>Ulrich Paquet</b>	Machine Learning for Recommender Systems and Online Advertising		9 & 10 January 2023
<b>Sebastian Bodenstein</b>	AlphaFold		9,10 & 11 January 2023
<b>Willie Brink</b>	Workshop: Introduction to deep Neural Networks		9, 10 & 11 January 2023
<b>Ferenc Huszár</b>	Theory of Deep Learning or Understanding Generalisation in Deep Learning this time around. I focus on questions of stochastic optimization in		11 & 12 January 2023

	overparametrized models these days, and I run a course on 'Theory of Deep Learning' in Cambridge.		
<b>Naila Murray</b>	Training deep vision models in low-data regimes.	Abstract: Large-scale vision models, pre-trained on vast quantities of often unlabelled data, give state-of-the-art performance when fine-tuned for a wide variety of downstream tasks. For these downstream tasks, the amount of training data available is usually orders of magnitude less than what was used during pre-training. And in many cases, there is little to no labelled data available. In this lecture series, I first present self-supervised approaches to visual model training. I will then discuss approaches to leveraging pre-trained models for vision tasks in low-data regimes. Focusing on video segmentation and object detection, I show that incorporating domain-specific and modality-specific inductive biases lead to improved model performance when training data is	11, 12 & 13 January 2023
<b>David Forsyth</b>	Computer vision for autonomous vehicles.		12 & 13 January 2023
<b>Sara Beery</b>	Domain adaptation and generalization from both a theoretical and an applied point of view.		12 & 13 January 2023
<b>Week 2</b>			
<b>NicuSebe</b>	1. Learning to Adapt: Adapting Deep Models to Domain and Semantic Shift.		16 & 17 January 2023

	2. Deep Generative Models for Image/Video Generation.		
	3. Playable Video Generation and Playable Environments		
<p><b>Xavier Alameda-Pineda</b></p>	<p>Unsupervised Probabilistic Learning with Latent Variables: from the Expectation-Maximisation Algorithm to Deep Variational Models.</p>	<p>Abstract: In this series of lectures, we will focus on probabilistic models and associated algorithms for unsupervised learning. We will first present the expectation-maximisation algorithm, which is a foundational tool in the field, allowing maximum likelihood estimation of the model's parameters as well as inference of the latent variables. The limitations of the EM algorithm, mainly the need for a closed-form posterior distribution, will motivate the need for approximate techniques, among which we can find variational inference. The coupling with deep learning, mainly via variational auto-encoders, will then appear as a natural crossing between probabilistic modeling and deep unsupervised learning. Applications as well as computational aspects of these models will also be discussed.</p>	<p>16 &amp; 17 January 2023</p>

<b>Inga Strümke</b>	XAI(Xeplainable Artificial Intelligence)		16 & 17 January 2023
<b>Matthias Bauer</b>	Probabilistic-generative modelling and would cover either all or some of the following: VAEs (with a bit of background on variational inference + potentially probabilistic PCA), normalizing flows, and potentially an outlook/connection to diffusion models.		18 & 19 January 2023
<b>Matt Jones</b>	Lectures 1: The Robots are Coming, be Afraid! How to fight back with Human-Centered Machine Learning Approaches	While many researchers and developers are excited about the possibilities of ML, everyday folk are often worried about the impact on their lives and livelihoods. We consider these concerns and use them to motivate the need for human-centered AI design and development approaches.	18 & 19 January 2023
	Lecture 2: Methodologies for Diversifying Perspectives in ML Design and Development.	ML systems have been criticised for biases caused by both the training data and model structures. In this second lecture, we discuss ways of involving richer perspectives from both the Global South and Global North in the development cycle.	

	3.Case Study: Speech Interactions designed and developed using human-centred machine learning approaches.	Speech Interactions designed and developed using human-centred machine learning approaches. In this last lecture, we bring everything together by showing how spoken language systems developed with Global South communities can not only provide new forms of service but provide challenges to conventional ML tools and techniques.	
<b>Steve Kroon</b>	Normalizing flows and VAEs.		19 & 20 January 2023
<b>Bettina Berendt</b>	AI Ethics: Some Important Concepts AI Ethics: Calling Bullshit AI Snake Oil		19 & 20 January 2023