New Learning Methods
Organizer: Nicu Sebe, University of Trento, Italy

The goal of this panel is to discuss new learning paradigms in AI from different perspectives covering online and continuous learning, efficient and cost-reduced approaches for a Green AI as well as learning in the real-world, i.e., dynamically learn to generalize and adapt across domains, labels and tasks simultaneously, and do so in a data-efficient fashion. The participating panelists have a varying level of research experience and are covering different areas of expertise going from computer vision to multimedia and machine learning.

Learning Compatible Representations via Stationarity
Alberto del Bimbo, University of Florence

Natural intelligent systems learn incrementally by continuously receiving information over time. They learn new concepts adapting to changes in the environment by leveraging past experiences. A remarkable capability of these systems is that learning of new concepts is achieved while *not* for getting previous ones. Current intelligent computer vision systems perform recognition exploiting internal feature representations learned by Deep Convolutional Neural Network models. In the case in which these systems receive novel data, representation models may require to be updated to improve their recognition capabilities. Such process can be computationally expensive or even infeasible. Previously acquired knowledge is likely to be forgotten. Learning compatible representation – i.e. such that the features before and after the learning update can be directly compared - has recently received increasing attention. We demonstrate that compatibility can be achieved by feature stationarity i.e. preallocating a special classifier with a large number of output nodes in which the weights are fixed (i.e., not undergoing learning) and set to values taken from the coordinate vertices of regular polytopes. The classification layer so defined allows learning stationary features that do not change their geometric configuration when novel classes are incorporated into the model.

Alberto Del Bimbo is Professor of Computer Engineering at the University of Firenze, Italy. He started his career in IBM and has been working in Computer Vision and Multimedia Analysis since 1988 when he joined the Academy. He is the author of over 300 scientific publications that have appeared in some of the most prestigious journals and conference proceedings. He has served as the Associate Editor of Pattern Recognition, Multimedia Tools and Applications, Pattern Analysis and Applications, IEEE Transactions on Pattern Analysis and Machine Intelligence and IEEE Transactions on Multimedia and is the Editor in Chief of ACM Transactions on Multimedia Computing, Communications, and Applications. He was the General Chair of ECCV 2012, the European Conference on Computer Vision, ACM ICMR 2011, the International Conference on Multimedia Retrieval, ACM MULTIMEDIA 2010, the International Conference on Multimedia and IEEE ICMCS 1999, the International Conference

This event is supported by the European Union's Horizon 2020 Research and Innovation programme under grant number 951911 – AI4Media

on Multimedia Computing and Systems. Presently, Alberto Del Bimbo is the Director of MICC, Media Integration and Communication Center at the University of Firenze, and the Coordinator of the NEMECH initiative, a Competence Center jointly established by Tuscany Regional Government and University of Firenze on New Media for Cultural Heritage.

New Forms of Learning for Green AI
Radu Tudor Ionescu, University of Bucharest

Over the past few years, handcrafted models for various machine learning tasks have been replaced by trainable models based on deep learning, due to the superior performance of the latter models. The main drivers for the performance improvements are the progressively deeper neural models and the increasingly larger data sets. However, since the 1980s, neural networks are still being trained with some variant of stochastic gradient descent, under the supervised learning paradigm. Successful convergence depends on carefully babysitting the training process by adjusting various parameters and on using large amounts of data. Hence, the training process of deep models is inefficient, even raising concerns about greenhouse gas emissions caused by large power consumptions. We believe that the next challenge in AI is to develop new forms of learning and more efficient neural models, taking a considerable leap towards Green AI. Solving this challenge is extremely important for reducing computational costs, thus (i) leading to significantly lower CO2 emissions caused by training and running the models and (ii) enabling new applications of deep neural models in certain scenarios (requiring real-time or embedded processing) which were not previously possible due to the high computational costs. These two benefits will have a positive impact in most areas where machine learning is currently used, e.g. autonomous driving, robotics, medicine, industry, agriculture, security and so on.

Radu Tudor Ionescu is Professor at the University of Bucharest, Romania. He completed his PhD at the University of Bucharest in 2013 and received the 2014 Award for Outstanding Doctoral Research in the field of Computer Science from the Romanian Ad Astra Association. His research interests include machine learning, computer vision, text mining, image processing and medical imaging. He published over 90 articles at international peer-reviewed conferences and journals, e.g. CVPR, ICCV, NeurIPS, ACL, EMNLP and TPAMI, and a research monograph with Springer. He received the "Caianiello Best Young Paper Award" at ICIAP 2013 and the "Danubius Young Scientist Award 2018 for Romania" from the Austrian Federal Ministry of Education, Science and Research.

This event is supported by the European Union's Horizon 2020 Research and Innovation programme under grant number 951911 – AI4Media

Real-World Learning
Cees Snoek, University of Amsterdam

Progress in artificial intelligence has been astonishing in the past decade. Cars self-driving on highways, machines beating go-masters, and cameras categorizing images in a pixel-precise fashion are now common place, thanks to data-and-label supervised deep learning. Despite the impressive advances, it is becoming increasingly clear that deep learning networks are heavily biased towards their training conditions and become brittle when deployed under real-world situations that differ from those perceived during learning in terms of data, labels and objectives. Simply scaling-up along all dimensions at training time seems a dead end, not only because of the compute, storage and ethical expenses, but especially as humans are easily able to generalize robustly in a data-efficient fashion. Several learning paradigms have been proposed to account for the limitations of deep learning with the i.i.d. assumption. Shifting data distributions are attacked by domain adaptation and domain generalization, changing label vocabularies are the topic of interest in zero-shot, open set and open world learning, while varying objectives are covered in meta-learning and continual learning regimes. However, there is as of yet no learning methodology that can dynamically learn to generalize and adapt across domains, labels and tasks simultaneously, and do so in a data-efficient fashion. This is the ambitious long-term goal of ‘real-world learning’. Streaming video provides a natural medium to study real-world learning, as it comes with an ever-changing amount of data, labels and objectives. Moreover, it’s multimodal, spatiotemporal, causal and semantic consistency and compositionality offers the means for efficient learning and adaptation to make the goal a reality.

Cees Snoek is a full professor in computer science at the University of Amsterdam, where he heads the Video & Image Sense lab. He is also a director of three public-private AI-research labs: QUVA lab with Qualcomm, Atlas lab with TomTom and AIM lab with the Inception Institute of Artificial Intelligence. At University spin-off Kepler Vision Technologies he acts as Chief Scientific Officer. He is also co-founder of the Netherlands Innovation Center for Artificial Intelligence. He was previously visiting scientist at Carnegie Mellon University and UC Berkeley, head of R&D at University spin-off Euvision Technologies and managing principal engineer at Qualcomm Research Europe. His research interests focus on making sense of video and images. He has published over 250 refereed journal and conference papers in computer vision, multimedia analysis and machine learning and serves on the editorial board of IEEE Transactions on Pattern Analysis and Machine Intelligence. Cees is recipient of an NWO Veni career award, a Fulbright Junior Scholarship, an NWO Vidi career award, and the Netherlands Prize for ICT Research. Together with his Ph.D. students and Post-docs he has won several best paper awards.

This event is supported by the European Union’s Horizon 2020 Research and Innovation programme under grant number 951911 – AI4Media