Continual Learning

Name(s) and Affiliation(s) of the Lecturer(s):

Chang D. YOO
School of Electrical Engineering, KAIST
Daejeon, Republic of Korea
cd_yoo@kaist.ac.kr
+82-42-350-3470

Haeyong KANG
School of Electrical Engineering, KAIST
Daejeon, Republic of Korea
haeyong.kang@kaist.ac.kr

Tutorial Background and Objectives:

In the following paragraph, please clearly specify the background and objectives of your tutorial. What is the topic? Why is it necessary and timely? What are the specific learning outcomes / skills that participants can expect to take away?

In the dynamic landscape of machine learning, deploying algorithms that can adapt to evolving trends and concept drifts is paramount. Continual learning, also known as lifelong learning, addresses this need by enabling models to learn continuously from new data while retaining knowledge from previous iterations. This tutorial will delve into the recent seminal work on continual learning, focusing on its necessity in a world where data and trends are constantly changing. Participants will be introduced to the best-performing continual learning methods, particularly those based on the hypothesis lottery, which hypothesizes that among the many parameters in a neural network, there exists a subnetwork that is particularly suited to a given task. Additionally, the tutorial will cover techniques for few-shot adaptation, equipping participants with the ability to quickly adapt models to new tasks with minimal data. By the end of this tutorial, participants will gain a comprehensive understanding of continual learning, its applications, and practical skills in implementing these methods in their own projects.

Target Audience:

In the following paragraph, please describe the target audience for your tutorial. If participants need to have specific prerequisites, please outline them clearly here.

This tutorial is designed for researchers, data scientists, and machine learning practitioners who are interested in advancing their knowledge of continual learning techniques. Participants should have a solid foundation in machine learning concepts and be comfortable with neural networks and deep learning frameworks. Familiarity with Python programming and experience with libraries such as TensorFlow or PyTorch will be beneficial. This tutorial is ideal for individuals who seek to enhance their models' adaptability and performance in real-world, ever-changing environments. Whether you are working in academia or industry, if you aim to stay at the forefront of machine learning advancements, this tutorial will provide you with the essential tools and insights.

About the Lecturers:

Chang D. YOO is a professor at the School of Electrical Engineering, KAIST, and an adjunct professor in the Department of Computer Science at KAIST. He is also a professor in the Robotics Program at KAIST. He earned his Ph.D. in Electrical Engineering from MIT, his M.S. in Electrical Engineering from Cornell University, and his B.S. in Engineering and Applied Science from Caltech. Dr. Yoo's research interests include machine learning, computer vision, and audio/speech processing. He has served as Dean of the Office of International Relations and the Office of Special Projects and Institutional Relations at KAIST. He is a member of Tau Beta Pi, Sigma Xi, ISCA, and a senior member of IEEE. Additionally, he serves as Vice President of the Institute of Electronics and Information Engineers (IEIE) and has held roles such as editorial board member of the Journal of Signal Processing Systems, tutorial chair for Interspeech 2004 and ACCV 2012, and area chair for ICANN 2011, ICCV 2019, ECCV 2020, and ECCV 2024. His awards include the 2020 Engineering School Technology Innovation Award and the Achievement Award from the Institute of Electronics Engineers of Korea in 2011 and 2013.

Haeyong KANG is a postdoctoral researcher at KAIST. He completed his Ph.D. with a dissertation titled "Forget-free Subnetworks for Life-long Learning" in August 2023. His research interests include sparse/pruned neural networks, ResNets, transformers, continual/incremental learning, multimodal representation learning, and imbalanced dataset learning. Specific topics of interest include task/class incremental learning, video incremental learning, continual pre-training of language models, and continual learning of knowledge graph embedding. He has reviewed for IEEE Letters/Transactions, Expert Systems With Applications, Springer, NeurIPS, ICML, ICLR, and CoLLAs.