AI studies: AI PhD Excellence. AI Curriculum. AI Science andEngineering?

Abstract: This panel has at least the following related discussion topics:

a) How can we 'quantitatively' define AI PhD excellence, so that young PhD graduates can enter a successful academic or industrial career?
b) What coursework should a quality AI MSc/PhD program contain?
c) Can we define a new 'AI Science and Engineering' discipline that is different from CS/CSE/ECE and be served at undergraduate level?

Panel Organizer: Prof. Ioannis Pitas (AUTH),

Panelists: Prof. A. Geiger (U Tuebingen), Prof. B. O'Sullivan (UCC), Prof. B. Schoelkopf, Prof.Ioannis Pitas (AUTH).

Statement of Prof. Ioannis Pitas on 'AI Science and Engineering'

In the past 2 decades the term 'Science' was coined all too frequently in Computer Science and Engineering (CSE) research, e.g., as Data Science or Web Science sometimes for no sound scientific reasons. Why is there such a recent proliferation of the term Science? Well, it certainly coincides with the era of competitive R&D funding. It has to do with the need to rejuvenate a discipline that gets oldish. Sometimes, it can be attributed just to vanity and a need to impress.

In light of the above, can we propose the term 'AI Science and Engineering' (AISE)? Right now, core AI is strongly related to Computer Science and Engineering. However, this relation is not rich enough for this field to progress and flourish. We badly need related research disciplines: Mathematics (lots of it), Systems Theory, Network Theory, Neurophysiology, Cognitive Neuroscience and Sociology, to name just a few. Then, we have major 'inexact' Sciences that come into play, ranging from Medicine to Law to Education to Politics. The relation of all these disciplines and AI is neither superficial nor coming out of an infinite greed of AI scientists. Indeed, mastering several of these disciplines is essential e.g., to mimic human intelligence or design ethical machines or design social e-interaction means that do not suffer from the current negative aspects of the social media.

Presently, AI education is served by MSc specializations and PhD studies in CS, CSE and Electrical/Computer Engineering (ECE) Departments. New AI BSc programs appear, primarily as a response to a demand for AI Scientists and Engineers. Will these domains ever converge to form a true new science to be taught at an undergraduate level as well? I cannot guarantee it, but I think that we do need such a confluence to address current research challenges, which well beyond current CSE and ECE curricula. Whether this science will be called AI Science and Engineering or Information Science and Engineering, this is something to be debated.

If I were to propose a definition of AISE, I would use the following: "AI Science and Engineering is the interdisciplinary, scientific study and engineering of Artificial Systems that mimic and/or surpass

human intelligence in information analysis and human interaction with the world". The involved core AISE disciplines are: Artificial Intelligence (AI), Machine Learning (ML), Robotics, Autonomous Systems, Digital Signal/Image processing and Analysis, Data Science and Data Analytics and Network Theory. They will be complemented with input from the following AISE- related disciplines from Liberal Sciences and Medicine: Cognitive Science, Neuroscience, Psychology, Philosophy, Linguistics and Sociology. Currently there is a real revolution in all these sciences, with respect to their applications in mind and social engineering. Recommendation systems is such a striking example. In this light, AI Science and Engineering studies can be complemented with a sister 'Mind and Social Science and Engineering' (MISSE) discipline that will expand the scope of the above-mentioned Liberal Sciences and Medicine.

Statement of Prof. Andreas Geiger on the 'AIDA BenchmarkPrinciples for Academic AI PhDs'

This statement has been compiled by Prof. Geiger and reflects opinions expressed by members of the following AIDA Committee. The work is in progress and will be finalized soon, under the coordination of Prof. Geiger.

Committee:

- Alyosha Efros <efros@eecs.berkeley.edu>
- Andreas Geiger <a.geiger@uni-tuebingen.de>
- Ioannis Pitas <pitas@csd.auth.gr>
- Mihaela van der Schaar <mv472@damtp.cam.ac.uk>
- Hector Geffner <hector.geffner@upf.edu>
- Bernhard Schölkopf <bs@tuebingen.mpg.de>

Goal:

- The goal of this committee is to establish principles and requirements for a PhD preparing for a successful academic career in AI.
- The benchmark principles shall help spread AI excellence across the world.
- They are general and apply to both, PhD studies conducted primarily in academia and PhD studies conducted partially or predominantly in industry. The authors of this document are convinced that there must not be any difference between a PhD conducted in academia and a PhD conducted in industry. Both should be equivalent.

Principles:

- **Research Excellence and Impact:** To establish the foundation for a successful academic career, PhD students should strive for research excellence and impact. As a rule of thumb, a successful PhD student should have authored at least 3 publications in top-tier conferences and/or journals as first author within 3 to 4 years. Top-tier conferences and journals in AI and ML related areas include
  - Core ML: NeurIPS, ICML, ICLR, UAI, AISTATS, COLT, AAAI, IJCAI, JMLR
  - Computer Vision: CVPR, ICCV, ECCV, IJCV, PAMI
  - Natural Language Processing: ACL, EMNLP, TACL

- Robotics: RSS, CoRL, AAMAS, ICAPS
- Data Mining: KDD, SIGIR, WSDM, KR

One publication could be in an equal-level top tier conference/journal in a different but AI-related area. Impact can be measured by the number of citations a publication receives within the first 3 or 5 years. Impactful work may be cited several hundred times within the first 3 years. However, citations vary across subfields, depending on their size. All research results including papers, source code and data should generally be given the possibility of publication as common in the field of AI and ML. Publication should not be hindered or significantly delayed by potential patenting processes.

- **Educational Background:** A successful academic career in AI necessitates the required theoretical and practical foundations acquired through coursework or in-depth self-study, during the overall course of (undergraduate and graduate) study. In this context, two areas of study are particularly relevant:
  - **Fundamentals:** Mathematics: Calculus, Linear Algebra, Probability Theory, Statistics, Optimization; Computer Science: Programming, Algorithms, Theoretical CS, Logic
  - **Artificial Intelligence:** Machine Learning Algorithms and Architectures (including Supervised Learning, Unsupervised Learning, Reinforcement Learning, Deep Learning), Statistical Learning Theory, Probabilistic Modeling, Causal and Symbolic AI, Computer Vision, NLP, Robotics, Planning, Constraint Reasoning.

  They can be complemented by courses in neighboring research areas, e.g., Data Science, Web Analytics, Neuroscience, Psychology (to name some examples) and/or application related courses, in particular courses on the societal aspects of AI. The above assumes that basic training in computing and programming has taken place during the undergraduate education.
- **Research Environment:** A PhD student should be closely supervised (at minimum monthly) by an expert in AI or ML research and be surrounded by peers (eg, PhD students and PostDocs) conducting research in AI or ML in order to acquire the necessary academic skill set. If the PhD studies are conducted at a lab where such an environment does not exist, eg, in industry or within another field where AI/ML techniques are applied but not developed, the PhD student should be co-supervised by an expert in AI or ML and conduct a research visit of at least 6 months during the beginning of the PhD period in a leading AI or ML research lab. In general, it is encouraged that PhD students collaborate or spend time with researchers from other countries to strengthen their network, eg, via PhD programs like the ELLIS PhD program. During the PhD studies, the PhD student should spend the majority of time on research related to the PhD and the minority of time on unrelated tasks (teaching, applications, industry work). For PhD studies conducted in industry, it is desirable to let the problem types be influenced by industry, but to assign maximal freedom to the PhD student in terms of conducting fundamental research in this area. Given the necessary compute required by many machine learning models, the PhD student should have access to a large enough (GPU) compute cluster during the entire period of study.

- **Continuous Learning:** A successful academic career requires the candidate to learn about recent advances and developments, for example through participating in short courses or attending conferences and workshops (at least one conference per year). During the beginning of the academic career, the PhD student should attend at least one summer school in machine learning or a related field, for example MLSS or ICVSS.

Andreas Geiger is a full professor at the University of Tübingen. Prior to this, he was a visiting professor at ETH Zürich and a group leader at the Max Planck Institute for Intelligent Systems. He studied at KIT, EPFL and MIT, and received his PhD degree in 2013 from the Karlsruhe Institute of Technology (KIT). His research interests are at the intersection of computer vision, machine learning and robotics, with a particular focus on 3D scene perception, deep representation learning, generative models and sensori-motor control in the context of autonomous driving. In 2012, he has published the KITTI vision benchmark suite which has become one of the most influential testbeds for evaluating stereo, optical flow, scene flow, detection, tracking, motion estimation and segmentation algorithms. His work has been recognized with several prizes, including the Everingham Prize, the IEEE PAMI Young Investigator Award, the Heinz Maier Leibnitz Prize of the German Science Foundation and the German Pattern Recognition Award. In 2013 and 2021 he received the CVPR best paper and best paper runner-up awards. He also received the best paper award at GCPR 2015 and 3DV 2015 as well as the best student paper award at 3DV 2017. In 2019, he was awarded a starting grant by the European Research Council. He is a board member of ELLIS and associate faculty of the International Max Planck Research School (IMPRS) for Intelligent Systems. He coordinates the ELLIS PhD and PostDoc program. He regularly serves as area chair and associate editor for several computer vision conferences and journals including CVPR, ICCV, ECCV, PAMI and IJCV, and is a program committee member of CVPR 2023.

Prof. Ioannis Pitas (IEEE fellow, IEEE Distinguished Lecturer, EURASIP fellow) received the Diploma and PhD degree in Electrical Engineering, both from the Aristotle University of Thessaloniki (AUTH), Greece. Since 1994, he has been a Professor at the Department of Informatics of AUTH and Director of the Artificial Intelligence and Information Analysis (AIIA) lab. He served as a Visiting Professor at several Universities.

His current interests are in the areas of computer vision, machine learning, autonomous systems, intelligent digital media, image/video processing, human-centred computing, affective computing, 3D imaging and biomedical imaging. He has published over 920 papers, contributed to 45 books in his areas of interest and edited or (co-)authored another 11 books. He has also been member of the program committee of many scientific conferences and workshops. In the past he served as Associate Editor or co-Editor of 13 international journals and General or Technical Chair of 5 international conferences. He delivered 98 keynote/invited speeches worldwide. He co-organized 33 conferences and participated in technical committees of 291 conferences. He participated in 71 R&D projects, primarily funded by the European Union and is/was principal investigator in 43 such projects. He is AUTH principal investigator in H2020 R&D projects Aerial Core, AI4Media (one of the 4 H2020 ICT48 AI flagship projects) and Horizon

Europe R&D projects AI4Europe, SIMAR. He is chair of the International AI Doctoral Academy (AIDA) https://www.i-aida.org/. He was chair and initiator of the IEEE Autonomous Systems Initiative https://ieeeasi.signalprocessingsociety.org/. Prof. Pitas lead the big European H2020 R&D project MULTIDRONE: https://multidrone.eu/ He has 34000+ citations to his work and h-index 86+. According to https://research.com/ he is ranked first in Greece and 319 worldwide in the field ofComputer Science (2022).