Educational renewal and educational development funds at Uppsala University

Computer assisted image analysis education in the era of deep learning and artificial intelligence – challenges and opportunities

Project report

Project leader: Nataša Sladoje, Department of Information Technology

1. Abstract

The subject of computerized image analysis has undergone dramatic changes in recent years. Modern deep learning and artificial intelligence have drastically advanced the state of the art, and have challenged traditional assertions about what is possible to achieve. Among students at Uppsala University, there is a strong and growing demand for courses in this field, combined with high expectations that the contents cover the latest progress in the field.

This project aimed at restructuring and modernizing the undergraduate courses in computerized image analysis, in response to the very rapid development of the subject. Aiming at offering a strong theoretical foundation, but also developing a set of skills towards targeting diverse applications, we have developed and implemented two new courses: 1MD110 *Introduction to Image Analysis* (10hp) and 1MD120 *Deep learning for Image Analysis* (7.5hp). They were given for the first time during the academic year 2020/21 and are included in several master's programmes at the Department of Information Technology.

The conducted evaluation indicated that we reached, and exceeded, the project goals. The developed courses showed to be highly appreciated by the students, while responding to all the stated aims. The students emphasize the relevance and timeliness of the included content, as well as very well selected teaching approaches to ensure student activation and learning.

2. Project description

Aim and goal

The use of digital images and visual data in society has exploded over the past decade. Demand for professionals with knowledge of image data analysis is ever increasing, as computerized image analysis is finding growing use within medical diagnostics, virtual surgery, life sciences, humanities, robotics, security, and many more. State-of-the-art image analysis strongly relies on, and drives the development of, the very latest tools of artificial intelligence and deep learning. The joint rapid advancement of both fields opens up new opportunities but also new challenges, both technical and ethical, which require new professional profiles and training tracks. We have observed a need to adapt the education in image analysis to accommodate for these changes.

The impact of deep learning methods in the field of image analysis is so large that the content related to them should not just be added on top of existing course topics, but instead be incorporated as a central concept. At the same time, the fundamentals of the field of (sometimes called "classic") image processing should not be neglected. It is continuously confirmed that deep learning based image analysis methods usually perform extremely well, but are considered non-intuitive and sometimes difficult to analyse and understand. They are typically referred to as "magic black boxes", a perception that may lead to their noncritical utilization.

The overall project goal has been to evaluate and conclude what is the best way to provide education in computer assisted image analysis for students of the 21st century; how to best combine education targeting the very latest techniques while not sacrificing fundamental theoretical knowledge. We aimed to reach this goal by taking a holistic view on the joint field of image analysis and machine learning, and evaluate different approaches to restructuring our current undergraduate courses (Computer assisted Image analysis I and II, 1TD396 and 1TD398) to reflect and respond to the recent rapid developments.

Our aim has been to enable students to get insight and acquire skills to evaluate advantages and disadvantages of both classic and deep learning approaches, and to critically approach powerful algorithms, understanding their main ideas and properties, and, based on that, be able to select, combine and confidently utilize the most suitable tools for a task at hand.

Project partners and their responsibilities

The project team consisted of four experienced lecturers: Nataša Sladoje, Filip Malmberg, Joakim Lindblad, and Robin Strand, all being actively involved in a number of image analysis courses at undergraduate and graduate level. The team is well aware of the current trends and needs in the field (as active researchers and with significant industry connections), as well as of the current content of the existing courses, and the way this content has developed over time.

We have, through active joint discussion, planned the restructuring of the existing courses, with focus on most suitable pedagogical approaches. N. Sladoje was responsible for managing the project activities and leading the discussions. Based on this preparatory phase, N. Sladoje has developed and coordinated a new course 1MD110 *Introduction to Image Analysis* (10hp) given in Periods 1&2, 2020; Joakim Lindblad has developed and coordinated a new course 1MD120 *Deep Learning for Image Analysis* (7.5hp), given in Period 4 2021, and Filip Malmberg and Robin Strand have been actively involved in teaching (at 1MD110).

We have worked in close collaboration with other teachers, to ensure seamless integration of all the related concepts towards a holistic management of image analysis and machine learning education at the IT Department. Particular efforts were made to well connect to, and integrate notions and concepts of related courses such as the Statistical Machine Learning course (1RT700) as well as the faculty wide PhD course on Deep Learning given in Period 4 2021 (TUFF supported, J. Lindblad is one of the teachers). Four PhD students at the research program Computerized Image Processing acted as TAs at the new 1MD110 and 1MD120 courses, and considerably contributed to their development. Furthermore, we have contacted (in particular through activities of the Swedish Society for Automated Image Analysis) several industrial partners active in the field of image analysis and machine learning, and discussed their specific views and needs, with an aim to respond to them by suitably integrating them in the course content.

3. Implementation and method

To reach the above stated goal, we have offered education on both model-based and learning-based approaches, discussing their properties, evaluating and comparing them theoretically and practically. It was observed that a joint treatment of model-based (“classic”) and deep learning based approaches requires substantial pedagogical/teaching effort, both due to the constant and rapid changes in state of the art, and also because such a path has not been taken previously and needed to be carefully thought through.

The project was implemented in three phases:

1. We have started with an extensive analysis of the existing situation (courses, content, and approaches) and planning of the changes;
2. Based on this analysis, we have developed and implemented two new master courses at the Department of Information Technology – 1MD110 *Introduction to Image Analysis*, with a focus on model-based approaches, and 1MD120 *Deep Learning for Image Analysis*, with a focus on learning-based approaches;
3. We have evaluated the courses by conducting surveys and following students' performance. Furthermore, we have assessed the position of the two courses in the wider context of education in the field of image analysis and machine learning.

We have focused on finding suitable content, structure and progression of the two courses, as well as an appropriate selection of teaching and assessment methods. Even though there is a high level of independence of the two courses, there is also a natural progression (including formal: 1MD110 is a suggested prerequisite for 1MD120) and synergy between them. Taken together, these courses offer a holistic overview on modern image analysis, and stimulate development of a wide range of knowledge and skills, due to use of diverse teaching methods, aligned with the teaching outcomes and the increasing complexity of the addressed content and tasks.

Structure and progression of the courses

An important part of the project was to find an optimal balance in the taught material, i.e., between the latest techniques providing impressive and thus inspirational performance and fundamental knowledge expected to stand the test of time. During the preparatory phase, we have explored and discussed:

- how to best combine and order the content to give it suitable focus and overall flow,
- how to suitably mix theory and practice,
- how to best relate model-based and learning-based approaches,
- how to best support the heterogeneous student group,
- how to best use online resources,
- how to assess and evaluate student performance while ensuring constructive alignment with novel course content and learning outcomes.

Lectures

To stimulate understanding of the methods and critical thinking, but also awareness of the rapid development of the field, we have decided to introduce the relevant concepts by analysing and comparing different approaches to address the same task, thereby enabling students to both learn about the algorithms, their basic ideas and implementations, and also evaluate them in different contexts, and based on different criteria. We have redesigned all lectures in the previously existing image analysis courses and adapted the content to reflect modern educational practices. The completely new deep learning-based content is developed with inspiration from successful courses all over the world. We have also performed a survey of available course literature and online resources. We have not found a unique text-book suitable for the envisioned contents, but rather had to combine different resources, including recent research publications; this well reflected the main challenge addressed in the project. We have discussed the feasibility of creating and publishing our own lecture notes/text book, to fill in this gap. The textbook "MACHINE LEARNING A First Course for Engineers and Scientists" (Lindholm, 2021), being written during the project time by colleagues at the IT dept., provided a good addition to the literature of the 1MD120 course.

Practical (computer) exercises

Practical exercises are a very important part of the image analysis and machine learning courses. They provide deeper exploration of topics learned and valuable hands-on experience leading to deeper understanding and improved learning. Our conducted evaluations of the previously existing

courses, and experience from other projects where courses have been redesigned to consider the impact of deep learning, such as e.g. (Wolfer, 2018), show that this format (practical exercises) is highly appreciated. We have decided to keep the format in the restructured courses, but also to thoughtfully and considerably redesign the exercises to match the new content, and to further support our pedagogical approach. The development of the new 1MD120 required creation of a significant amount of new material in this direction, which included programming tasks, group assignments and also a Kaggle competition on a relevant problem. Our general aim was to ensure that students develop a deeper understanding of the methods, in particular emphasizing and utilizing method evaluation techniques and discussion of the obtained results (in 1MD110 typically during seminar sessions), but also to provide the students practical experience with modern environments for image analysis and machine learning.

Assessment (examination)

The workload of assessing the student reports and exams is rather high, and increases due to the increasing number of students interested in the field. One of the goals of this project was to investigate examination forms that maintain the formative and student active aspects of the oral examination, combined with the advantages of written reports, while also being sustainable with limited teaching staff and growing student groups.

4. Theoretical foundation

We have strived for increased active learning within the courses through activating forms of teaching (Pedagogiskt program UU, 2018) where in particular the critical evaluation of methods of different types will support reflection. Good timing between delivered material, lectures, and exercises enables increased usage of just-in-time teaching. Students practice critical thinking and evaluation strategies through analysis and assessment of the learned content on practical tasks. A range of skills (presentation, oral and written, team-working, creativity) are further developed in suitable group projects and practical exercises. Practical exercises and projects have successfully been used in other projects where courses have been redesigned to consider the impact of deep learning/AI in image processing (Wolfer, 2018) and embedded systems (Ergezer, 2018).
Our aim was to provide theoretical understanding and hands-on experience in a good balance as well as constructive alignment between teaching, learning, and assessment, where reasoning derived from well designed intended learning outcomes is driving the flow of the course contents (Biggs, 2011).

5. Results

Our initial project plan was to prepare detailed syllabi for the restructured courses, but to only implement a limited number of pilot-lectures in one of the existing image analysis courses, to evaluate the most relevant ideas of the new structure. The final outcome of the project highly exceeds this initial plan.

We have developed and implemented two new complete courses: *1MD110 Introduction to Image Analysis* (10hp), given in Period 1&2, and *1MD120 Deep learning for Image Analysis (7.5hp),* given in Period 4. Both are core courses of the Master's programme in Image Analysis and Machine Learning (offered at the Department of Information Technology), and are also offered to students of other master programmes at the IT Department, Uppsala University.

Based on the results of our initial analysis and discussion with colleagues, we concluded that to integrate and, in parallel, teach both model-based and learning-based methods, would possibly lead to a lack of focus on either approach. Although such an integrated teaching would simplify direct comparison and drawing of relations between the different methods, we concluded that such

a breadth-first approach would possibly miss the goal by trying to achieve too much at the same time. Instead it was concluded that a sequential approach, first presenting model-based techniques and only introducing learning-based approaches with few examples, followed by a separate course focused on deep learning techniques from the bottom up, would provide a better pacing of the material and allow delving deeper into relevant topics without losing focus. Such a separated approach also fits well within the landscape of other machine learning courses at the IT department, where students have an opportunity to, in particular, strengthen their machine learning knowledge before progressing to the 1MD120 course.

1MD110 Introduction to Image Analysis includes a broad range of topics covered in 17 lectures. Keeping in mind that the students taking this course have a deeper interest in the field, our focus was on providing a stable foundation for their further work. We have addressed a variety of topics, including terminology and main concepts, but all the time making sure to take a holistic approach, giving numerous examples from real life scenarios where the methods we presented can successfully be applied. Furthermore, the importance of scientific evaluation of the methods was emphasized from the start and throughout the course, and was practiced in all the computer labs and the mini-project.
We have addressed basic concepts of image representation and statistical analysis, various filtering techniques in spatial and frequency domain, main concepts of image segmentation, feature extraction and description, spatial alignment and motion analysis, classification and interpretation, as well as more theoretical concepts of discrete geometry, mathematical morphology, inverse problems and optimization, as tools for solving numerous image analysis tasks. Discussion on applicability and suitable evaluation strategies followed each set of the introduced methods. All the concepts were further explored, analysed, and discussed within 5 computer exercises, 2 seminar sessions, and one (mini-)project, with the content of increasing complexity, stimulating independent and creative problem-solving driven thinking and learning. Finally, one lecture was devoted to presentations by representatives from the industry, where relevance and applicability of the methods learned during the course were highlighted, in turn highlighting also the great need for an, in these topics, well educated workforce.

1MD120 Deep Learning for Image Analysis covers the whole span of deep neural network learning, from writing own regression model, via multilayer neural networks to state of the art transformer models. This is covered in 12 lectures, 3 individual assignments, 1 group assignment and 1 project challenge, of which the final two are presented and discussed in seminar sessions. The teaching in particular aimed to not progress too fast but ensure that the students feel comfortable with their knowledge about the more fundamental concepts. The initial 3 programming assignments covered this (fundamental) part in detail, building up a solid knowledge while leading towards the more advanced levels. The exam (3hp) also focused primarily on these fundamental parts of the course topics. A sequence of more advanced lectures, including two given by invited researchers actively working in the field, brought us towards inspiring state of the art techniques. The possibilities and challenges of different approaches were practically explored through group projects as well as an organized Kaggle deep learning challenge on Cancer cell classification, at the same time bringing the inspiration from a competition and also highlighting possible usage of the learnt methods within society (in this case healthcare). A final lecture of the course was devoted to fairness, accountability, transparency and ethics in AI, highlighting possible implications of biased (e.g. racial and gender) usage of these powerful techniques in society and what can be done to counteract such effects.

In both courses we have designed a whole range of practical exercises, introducing topics of different and increasing levels of complexity and requiring increasing levels of creativity, leading the students from very simple and guided tasks, all the way to their own originally designed solutions to complete image analysis and machine learning problems. Practical work was often

performed in teams (of 2-4 members, for different tasks in the two courses), stimulating active student learning, where students could discuss their solutions with both the TAs and their peers. Presentation of the results included both written reports, and oral presentations (depending on the task). By this we have stimulated development of students' writing (presentation) skills and deeper and more careful analysis of the tasks and solutions by each student individually. Oral presentations further contributed to the soft skill development and enabled experience of evaluation of own work in direct discussions with peers.

Being faced with the Covid-19 pandemic, with an effect and extent not envisioned during the time of project/course planning, we were forced to adjust, while developing the course, both the material and the teaching approach. The outcome is a highly valuable recorded material as well as other online content and tools (such as quizzes, webinars, and similar), covering two complete newly developed courses. We plan to extensively use this material in the future course editions, possibly supporting a flipped-classroom approach, but also further exploring and developing other possibilities. An observation from the course evaluations is that our approach to online teaching found good appreciation among the students, including explicits comments on, e.g., ease of keeping online presentations of group assignments for colleagues, or usefulness of online discussion forums for Q&A interaction with TAs and discussions with classmates.

Techniques for sustainable assessment of student reports and exams, envisioned to be explored within the project, were analysed with respect to restrictions imposed by the pandemic. The examinations were conducted online, and have imposed open-book exams (home-examination); this was not initially planned for any of the two developed courses. As a result, exam questions and tasks were (forced to be) suitably, and significantly modified. One direct observation is the increased difficulty of preventing disallowed techniques, such as student collaboration during the exam, for online examinations. We have considered diverse examination forms, including quizzes, projects, and written analysis and solution of the practical problems. Whereas grading in Studium offered increased comfort and efficiency, there is still space for improvement. We will further explore the possibilities with an aim to identify and implement efficient and sustainable, but also diverse and transparent assessment methods.

6. Evaluation of the project

Evaluation strategies

We have evaluated the project in several steps, by conducting surveys and discussions including different target groups. In the planning phase, we organized discussions with PhD students at the research program in Computerized Image Processing, since they are naturally well informed about state-of-the-art approaches in the field, and the needs for different skills and knowledge. In particular, four of these PhD students acted as TAs supporting the two courses. Furthermore, we kept frequent discussions with a number of lecturers at these two courses, to formatively assess changes and novel approaches; they were in a position to contribute to the project by directly introducing their own innovative ideas into their own lectures in the new courses. Finally, we have discussed the content of the courses with colleagues teaching image analysis and machine learning at the Department of Information Technology, in particular within courses included in (or planned to be developed for) the Master's Programme in Image Analysis and Machine Learning. We have invested efforts to ensure a holistic treatment of the taught content, enabling building of a solid foundation, as well as delivering suitable progression.

Further, we have conducted formative assessment during the course implementation time. This was of particular relevance, considering the specific conditions during the pandemic. Students gave highly valuable comments and suggestions, which gave good insight into the challenges they were facing, primarily due to considerably reduced interactions caused by online teaching. We

have dynamically reacted with suitable modifications which were appreciated. We have worked on utilizing the tools available in online teaching to our advantage, to the extent possible at courses not primarily designed and meant to be given online.

The final evaluation of the project is based on the final assessment of the two developed courses, conducted at their ends and targeting all the participating students.

From the evaluation of 1MD110 Introduction to Image Analysis:

Overall satisfaction with the course is graded 4 (median) and 4.3 (mean), with 5 being maximum.
Overall relevance of the course is graded 5 (median) and 4.5 (mean), with 5 being maximum.
Overall, active participation of the students is self-assessed as 4 (median), and 3.6 (mean), where 5 indicates a very high level of activity. (Considering the challenges imposed by the pandemic, we see this as a rather high level of activity.)

Students appreciated all the teaching methods used in the course: lectures (4.2), computer labs (4.3), seminars (4.5) and mini-projects (4.8).
(The numbers indicate the mean grade of the relevance and contribution to learning of the different used methods, 5 being the maximum.)

The students appreciated the content (topics) covered during the course. The holistic approach to solving image analysis tasks was particularly trained through the second seminar session and the mini-project. Students clearly appreciated both, finding them creative and interesting. They also expressed satisfaction with the fact that the methods they learned about were relevant in industry, which was clear after the presentations given by two company representatives during the special guest lecture.

Some of the quotes from the survey, as a response to *"indicate the strengths of the course"*, are:

- *"Every concept is explained in detail in lectures. Assignments contributed highly to gain practical knowledge."*

In particular

- *"Mini-project, being a complete image analysis pipeline to solve on our own, completed with a report and presentation."*

Or, in summary:

- *"Basically everything, really enjoyed the course."*

From the evaluation of 1MD120 Deep Learning for Image Analysis:

Overall satisfaction: 94% of the students are highly, or completely satisfied with the course.
Overall relevance of the course is graded 4.8 (mean), with 5 being maximum.

Assignments and project work were seen as most contributing to learning, with lectures, seminars and student presentations not far behind, and all much appreciated. Furthermore, students appreciated competitions (*"competitions always end up being enjoyable, it seems").*This indicates that we manage in our stated goal to find a suitable balance in the taught material, as well as in the teaching methods.

Regarding the project aim to encourage and assist the students to take an active approach to learning, (their own and their peers), implemented through group assignments, presentations and seminars, we have observed that:

- On the statement "I took a great deal of responsibility for my own learning during the course", 88% of the students answered "Agree to a high extent" or "Agree completely".

- On the statement “I contributed to other students’ learning during the course”, 60% answered “Agree to a high extent” or “Agree completely”.

Particular statements related to seminars and oral presentations: *“They certainly got me motivated to [...] try new ideas*” and *“mycket givande att lyssna på de andra gruppernas presentationer*”. We take these responses as indications of the successful implementation of our project.

We have addressed the aim of the project to include state-of-the-art research in the course by inviting lecturers who delivered appreciated presentations of contemporary methods and techniques, including, e.g., less than a week old material(!) from Facebook AI on self-supervised learning with vision transformers (Caron, 2021). Teaching can hardly become more up to date than that. Students clearly appreciated such a broad range (*“Very nice that we went from the basics all the way to state of the art.”).* While our aim to respond to the needs of industry and requests from students for state-of-the-art content was clearly met, we will evaluate modifications to the taken approach, considering an observation by one of the students that the most up-to-date content might better fit in a separate course, or – as also suggested – in a larger (and further extended) course.

Finally, students appreciated the conducted examination, which was found to be well aligned with the teaching methods and learning outcomes, targeting understanding of the core concepts and their applicability, and reasoning about their suitability and failure modes, towards their confident generalization and usage in practice.

7. Reflection and discussion

The overall goal of the project was to evaluate and conclude what is the best way to provide education in computer assisted image analysis for students of the 21st century. In particular, we wanted to identify approaches to best combine education targeting the very latest published techniques, while not sacrificing fundamental theoretical knowledge. In that sense, the project aimed at contributing with new ideas and new forms related to:

- promotion of a modern view of image analysis;
- education which supports and encourages critical thinking w.r.t artificial intelligence, promoting understanding and evaluation (including ethical aspects) of the powerful algorithms and their suitable selection and use in a variety of application fields;
- modernization of course literature and effective usage of online resources;
- evaluation of novel approaches to assessment of computer exercises, and learning outcomes in general.

We have reached, and exceeded, all the project goals. Highly relevant for the dynamics of project implementation was the Covid-19 pandemics, which imposed online education and, with that, numerous unexpected challenges. We planned to develop and (in a longer run) implement two courses delivered in the classrooms, assuming physical interactions with and between students. Instead, we were forced to adjust the plans (content and teaching methods) to online teaching.

A particular challenge was to stimulate students’ interaction. In particular, seminar sessions were a highly appreciated teaching method towards interactive and active student learning. To stimulate active and rewarding teamwork between students who have never met each other physically required special attention. In the 1MD120 course, discussion and interactions were supported via topic- and task-specific threads on a suitable discussion platform (Studium), which also became the main channel of reaching out to TAs for assistance and feedback on the assignments, something which worked very well. These more formal communication channels in Studium were (based on student suggestion) complemented with a less formal Slack channel, which very

successfully turned out to become a much appreciated and active forum for student-to-student interaction and discussions. This separation of communication channels, into one more and one less formal, is definitely something which we will use again in the future. Considering the circumstances imposed by the pandemic, we find the (self-assessed) level of students' activity highly satisfactory (far beyond our expectations).

We have adjusted the approaches taken in lecturing, but also during the lab sessions. A particularly valuable result of this adjustment (and of this project) is the produced lecturing material, which includes pre-recorded video-lectures for the two complete courses, often supported by thematic quizzes for assessment of understanding of the presented content. This material will, in the future, offer possibilities for twisted classroom teaching approaches (Berrett, 2012), but also exploration of additional and innovative teaching approaches which can rely on it.

We have, based on conducted formative assessment, learned to utilize opportunities of online teaching: examples include introduction of more frequent and better distributed drop-in help sessions for computer exercises, instead of the usual scheduled 4-hour on-campus sessions, but also appreciating the comfort of online oral presentations, typically in parallel break-out rooms. We will further evaluate the potential of online tools, to best incorporate them in our teaching also after going back to the classrooms.

Finally, examinations conducted in a form of open-book home exams imposed, but also stimulated, new design of exam tasks/questions, with increased focus on assessment of understanding, synthesis, and evaluation, promoting constructive alignment. We will continue along this direction, further improving sustainability in examination, while strengthening the taken approach, also highly appreciated by students.

8. Conclusions and vision

Based on the performed evaluation we conclude that the project has been very successful, hughly exceeding the initial plan and established goals. The modernized courses have an improved structure, which has created a clearly positive effect on the student performance, and is also highly appreciated as indicated in the received feedback. The renewed structure will therefore be kept, as long as it stays fit for the dynamic and rapidly changing field of education. Our continued communication with the industrial partners, and our active collaboration with researchers in the field, will provide insight in the most relevant directions to be taken to ensure sustainable development.

Furthermore, the restructuring initiated within this project has further positive influence on two other new advanced courses that will build on 1MD110 and 1MD120: (i) *Software Development Project in Image Analysis and Machine Learning*, and (ii) *Advanced Image Analysis*. Both are offered in year 2 of the Master's programme in Image Analysis and Machine Learning and will be given first time in the Autumn 2021. These two courses will further deepen students' understanding of the field and will put in use modern powerful tools available to solve a wide range of practical image analysis problems.

Both image analysis and machine learning are highly interdisciplinary fields, with a wide range of applications. By restructuring the courses, and suitably integrating relevant data-driven machine/deep learning approaches into the image analysis education, we have further enabled close collaboration and synergy around machine learning and artificial intelligence initiatives at the IT department, but also opened possibilities for enhanced collaboration with other both education- and research-focused groups relying on image data analysis such as material science, medical and biomedical ones. The modernized courses are very attractive for interdisciplinary and

intersectoral (industrial) collaborators, and we expect increased interest for both the courses, but also collaborative connections.

9. Reporting

We find not only the main result of the project – two new developed courses which address rapidly developing research and application fields by carefully selected range of teaching methods – interesting for dissemination, but also the not expected (but accepted) challenges of adjustment to online education during pandemic. Both can (and already have), in our opinion, stimulate a number of relevant discussion topics with colleagues, leading to further improved use and development. We have, so far, not decided on a most suitable venue.

10. References

- Berrett, D. “How ‘flipping’ the classroom can improve the traditional lecture”. The chronicle of higher education, 12(19), pp.1-3, 2012.
- Biggs, J and Tang, C. “Teaching for Quality Learning at University”, McGraw-Hill and Open University Press, Maidenhead, 2011.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. “Emerging properties in self-supervised vision transformers”. *arXiv preprint arXiv:2104.14294*, 2021
- Ergezer, Mehmet. "Work in Progress: Designing Laboratory Work for a Novel Embedded AI Course." 2018 ASEE Annual Conference & Exposition
- Lindholm, A., N. Wahlström, F. Lindsten, and T.B. Schön Machine Learning - A First Course for Engineers and Scientists, Cambridge University Press, 2021. To appear.
- “Pedagogiskt program för Uppsala universitet”, 2018. http://regler.uu.se/digitalAssets/14/c_14251-l_3-k_pedagogisk-program-beslutat-2018-04-24.pdf
- Wolfer, James. "Implementing a Project-Based, Applied Deep Learning Course." *International Conference on Interactive Collaborative Learning*. Springer, Cham, 2018.