Designing an Engaging Curriculum for Advanced Topics in Deep Learning: A Pedagogical Development Project

Mostafa Mehdipour Ghazi

Pioneer Centre for Artificial Intelligence
Department of Computer Science
University of Copenhagen

Introduction

The rapid advancement of artificial intelligence (AI) and deep learning has created an urgent need for educational programs that stay aligned with emerging research, preparing students for advanced academic and professional work. As deep learning gains significance across academia and industry, it offers a unique opportunity to shape the next generation of researchers and professionals. This calls for the development of an effective curriculum that integrates well-established pedagogical frameworks to enhance student engagement, collaboration, and problem-solving skills.

This article outlines a project aimed at designing and implementing a new curriculum for the MSc in Computer Science program at the University of Copenhagen. The focus is on the course "Advanced Topics in Deep Learning" (2024/2025), which explores state-of-the-art deep learning techniques, including deep generative models, multimodal foundation models, and graph neural networks. By addressing cutting-edge topics, the course equips students with the skills and knowledge needed to contribute meaningfully to the field. Additionally, the project enables the instructor to stay updated on the latest advancements while enhancing teaching and mentoring skills.

The project has three main goals: first, to develop a comprehensive and up-to-date curriculum on advanced deep learning topics; second, to engage students with the latest advancements through active learning and effective teaching strategies; and third, to foster a strong research culture through interdisciplinary collaboration and student mentorship, preparing them for impactful master's theses and future research. The project incorporates feedback from student surveys, course evaluations, and faculty reflections from the initial course iteration, recognizing strengths such as content depth and interactive discussions while identifying areas for improvement, including course description clarity, assignment structure, and exercise session efficiency.

Background

Effective curriculum design in deep learning education requires a balance between theoretical foundations and hands-on application. Research highlights the importance of active learning strategies in fostering deeper engagement and knowledge retention (Freeman et al., 2014). Methods such as problem-based learning (Hmelo-Silver, 2004) and project-based teaching (Blumenfeld et al., 1991) have been shown to enhance student motivation and conceptual understanding by encouraging self-directed inquiry and real-world problem-solving. The Theory of Didactical Situations (TDS) (Brousseau, 2006) further supports this approach by structuring learning environments where students explore and construct knowledge collaboratively, shifting control dynamically between instructor and learner. In deep learning courses, collaborative assignments and peer discussions have been particularly effective in reinforcing complex concepts.

Another critical challenge in deep learning education, particularly in large-class settings, is maintaining engagement and providing effective feedback. Studies indicate that interactive teaching approaches, including structured discussions, real-time coding exercises, and peer evaluations, improve student participation and learning outcomes (Prince, 2004). However, large auditoriums and high student-to-instructor ratios pose difficulties in ensuring personalized feedback and engagement (Deslauriers et al., 2019). Digital tools and AI-powered coding assistants can help bridge this gap by providing scalable support and formative assessment opportunities (Luckin et al., 2016). Research suggests that formative assessments, including peer reviews, in-class discussions, and scaffolded assignments, contribute to deeper engagement and skill development (Black & Wiliam, 1998). Additionally, continuous assessment methods, such as iterative feedback loops through assignments and presentations, have been found to improve learning efficacy in technical subjects (Nicol et al., 2006).

Effective course implementation also depends on clear expectations, structured guidance, and interdisciplinary collaboration. Literature on curriculum alignment emphasizes that defining clear learning objectives and assessment criteria enhances student comprehension and course coherence (Biggs, 1996). Moreover, research mentorship and interdisciplinary projects foster a strong research culture, preparing students for independent inquiry and graduate-level research (Sadler, 2014). Addressing student diversity in academic backgrounds, particularly in fields requiring mathematical and computational rigor, necessitates pre-course assessments and flexible instructional strategies (Chi et al., 1989). Future iterations of deep learning courses should continue to integrate these pedagogical best practices while leveraging student feedback for iterative course refinement.

Activities

The project includes key activities aimed at achieving its goals, such as curriculum development, active learning, and research mentorship. The primary focus is on designing a syllabus that balances theoretical knowledge with practical implementation through hands-on exercises and research articles. The curriculum covers advanced deep learning topics, including deep generative models, multimodal foundation models, and graph neural networks.

The project applies teaching methods, for example, based on TDS and the principle of congruence. Interactive lectures, group discussions, and problem-based learning foster a collaborative and engaging classroom environment. Students are encouraged to conduct in-depth research through guided assignments and project work, preparing them for future research and thesis projects.

Curriculum and Structure

The newly developed Advanced Topics in Deep Learning (ATDL) course (https://kurser.ku.dk/course/ndak24003u/) replaced the previous Advanced Deep Learning (ADL) course to address gaps in cutting-edge content. Covering topics such as deep generative models and multimodal AI, it required a strong background in deep learning, programming, linear algebra, mathematical analysis, probability theory, and data modeling. The course workload included 28 hours of lectures, 14 hours of class instruction, 70 hours of preparation, and 94 hours of exercises. Students had access to AI-based tools like GitHub Copilot and prepaid GPU resources for assignments.

The course was divided into four major topics, each taught by different instructors based on their expertise. It aimed to provide in-depth insights into advanced deep learning methods, covering algorithms, theory, and tools in this rapidly evolving field. The structure included two weekly lectures (1×45 minutes and 2×45 minutes) and three exercise and presentation sessions facilitated by senior PhD students serving as teaching assistants (TAs). These sessions (3×45 minutes) promoted collaborative group work and peer learning.

Organization and Management

The course was managed through regular weekly meetings among faculty and TAs to address ongoing issues and feedback. These meetings covered student concerns raised via the teaching Absalon platform (https://absalon.instructure.com/courses/76879) or during lectures, adjustments to uncovered topics, and suggestions for improving course content and delivery in future iterations. Key discussion points included assignment design and grading to ensure consistent and constructive feedback across instructors. The course coordinator handled any outstanding logistical and pedagogical concerns within the teaching team.

While mathematical proofs and traditional lectures remained essential, parts of the lectures incorporated TDS and active learning to enhance engagement and comprehension. For example, deep generative models were introduced through real-world applications like image generation in healthcare (devolution). Students then completed a small exercise, working in groups before formal theory was presented (activation). After posing a challenging question, they discussed their approaches in a safe environment (formulation). Feedback guided them toward optimal solutions while expanding on peer discussions (validation). Finally, class discussions were summarized with established literature to contextualize their efforts (institutionalization). This fostered dynamic interaction and balanced teacher-student control.

Assessment and Feedback

Before engaging in activities, students received clear instructions and learning outcomes for each lecture. Feedback was provided through class discussions, where students were encouraged to clarify their understanding and ask questions about upcoming tasks, real-time input during exercises and selected paper presentations, and detailed qualitative feedback on assignments after their submissions. In addition, Absalon served as a platform for students to ask questions about teaching materials and assignments. The formative peer feedback was instrumental in helping students self-correct while engaging with the material and collaborating in groups.

The assessment strategy was continuous, requiring three written assignments to pass. The final grade was based on an overall assessment of these assignments, with no individual grades but qualitative feedback provided. The first assignment involved reviewing and summarizing the synopsis of a selected research paper and submitting a report. The second needed implementing a baseline method from the chosen paper or reproducing one of its main results, along with a report and a link to the implementations or results. The final assignment was to write a short scientific paper with additional experiments related to the selected paper. The first assignment was intentionally simpler, helping students adapt to the feedback process and improve in later tasks. To promote engagement and discussion across topics, students were required to read and present one paper per group, per topic, after the relevant lecture. However, they were free to choose any topic for their final assignment and presentation.

Results

The project's success was evaluated through formal course evaluations, direct student feedback (both oral and written), and assessment performance analysis from assignments. Students' interest in pursuing research projects and master's theses in the field during and after the course completed also indicated its impact. The ATDL course was assessed through presentations and homework assignments for over 60 students with 7.5 ECTS.

Overall, students provided positive feedback, with course evaluations averaging above 4 out of 5. They praised the instructor's passion and the depth of content while suggesting improvements in lecture organization and delivery. Table 1 summarizes student ratings for the course at the University of Copenhagen (https://evaluering.ku.dk/), highlighting feedback on key aspects such as clarity, engagement, and instructional effectiveness.

Table 1. Summary of student ratings for teaching Advanced Topics in Deep Learning (ATDL) course at the University of Copenhagen in 2024.

	Evaluation Metric *5-scale Rating (Mean ± SD)*
Course Content Communication *Clarity and Accuracy*	4.10 ± 1.10
Learning Outcomes *Interest and Engagement*	4.40 ± 0.72
Instruction Language *Clarity and Expression*	4.40 ± 0.73

Student feedback highlighted both strengths and areas for improvement. For example, while the paper presentation format was well-received, students requested deeper discussions and better access to presentation materials. Improving exercise session efficiency and clarifying assignment expectations were also noted as priorities.

Discussion

Given that the ATDL course was newly developed, some inconsistencies arose in lecture delivery. While weekly meetings helped reduce content overlap, this remained an area for improvement. Additionally, structural challenges emerged as the course replaced the previous ADL course (offered in Block 4). The new course was scheduled in Block 1, while its prerequisite, the renamed Deep Learning (DL) course, being offered in Block 2. As a result, some students took ATDL without completing the prerequisite, requiring foundational material to be covered, while others had already mastered these concepts in the ADL course. This potentially reduced time for more advanced content. Regular reviews of course content and feedback from students and lecturers would help ensure that all material is appropriately challenging and relevant.

The course was well-organized and aligned with its intended learning outcomes. Regular coordination meetings ensured smooth progress despite multiple instructors and assistants. A clear task division among instructors ensured consistency across topics, which was crucial for a complex course. However, clearer communication of expectations and management refinements could improve alignment. Some students were uncertain about focus areas and the balance between lectures and exercises. Revising the course description to clarify student engagement, including time allocation for group work and presentations, would enhance their understanding of learning outcomes and course structure.

The feedback and assessment strategies provided multiple opportunities for both formative and summative feedback. Before each lecture, course materials and references were shared on Absalon, allowing students to prepare in advance. After the lecture, slides and a list of 10 selected papers were provided, enabling students to choose a paper of interest for their group assignments. This structure ensured that feedback and assessment were integrated throughout the learning process. However, incorporating more formative assessments, such as small quizzes or reflective exercises, could give students clearer insights into their progress before major assignments. Future iterations should also refine assignment descriptions, improve feedback consistency through clearer grading rubrics and timely, detailed feedback, and support diverse learning needs with additional resources, personalized guidance, and flexible assignment formats. Emphasizing advanced topics in denoising diffusion models and increasing hands-on programming exercises will further enhance the learning experience and better align the course with student expectations.

The diversity in students' academic backgrounds and familiarity with deep learning and mathematical concepts, while balancing theoretical and practical aspects, posed other challenges. For instance, during the lecture on implicit deep generative models, some students lacked background in differential equations, which was not part of their bachelor's or master's curriculum. Future iterations should refine the course description to set clearer expectations, explicitly outline prerequisites, and incorporate pre-course assessments to better tailor the learning experience.

To partially address the diversity issue, students were grouped into teams of 3-4 to work on papers of interest, allowing flexibility in learning and peer support. However, grouping posed challenges, leaving some students feeling isolated. In such cases, we assisted with group formation and ensured transparency, so students understood the process. Additionally, monitoring individual contributions within group work was difficult. To address this, students provided self-reports on their contributions and participated in group presentations. Nevertheless, more attention is needed to ensure all students, regardless of group dynamics, receive equal support and remain engaged in the learning process.

Conclusion

This project marked a significant step in enhancing the ATDL course by aligning curriculum design with pedagogical best practices and emerging research. Through active learning, research mentorship, and continuous feedback, the course aimed to prepare students for impactful academic and professional careers in AI. Future iterations will incorporate student feedback to refine and expand the course, ensuring its continued relevance and effectiveness in an evolving field.

References

Freeman, S., Eddy, S. L., McDonough, M., Smith, M. K., Okoroafor, N., Jordt, H., & Wenderoth, M. P. (2014). Active learning increases student performance in science, engineering, and mathematics. Proceedings of the National Academy of Sciences, 111(23), 8410.

Hmelo-Silver, C. E. (2004). Problem-based learning: What and how do students learn? Educational Psychology Review, 16(3), 235-266.

Blumenfeld, P. C., Soloway, E., Marx, R. W., Krajcik, J. S., Guzdial, M., & Palincsar, A. (1991). Motivating project-based learning: Sustaining the doing, supporting the learning. Educational Psychologist, 26(3-4), 369-398.

Brousseau, G. (2006). Theory of didactical situations in mathematics: Didactique des mathématiques, 1970-1990 (Vol. 19). Springer Science & Business Media.

Prince, M. (2004). Does active learning work? A review of the research. Journal of Engineering Education, 93(3), 223-231.

Deslauriers, L., McCarty, L. S., Miller, K., Callaghan, K., & Kestin, G. (2019). Measuring actual learning versus feeling of learning in response to being actively engaged in the classroom. Proceedings of the National Academy of Sciences, 116(39), 19251-19257.

Luckin, R., & Holmes, W. (2016). Intelligence unleashed: An argument for AI in education. Pearson Education.

Black, P., & Wiliam, D. (1998). Assessment and classroom learning assessment in education. Principles, Policy & Practice, 5 (1), 7-74.

Nicol, D. J., & Macfarlane-Dick, D. (2006). Formative assessment and self-regulated learning: A model and seven principles of good feedback practice. Studies in Higher Education, 31(2), 199-218.

Biggs, J. (1996). Enhancing teaching through constructive alignment. Higher Education, 32(3), 347-364.

Sadler, D. R. (2014). Beyond feedback: Developing student capability in complex appraisal. In Approaches to Assessment that Enhance Learning in Higher Education (pp. 45-60). Routledge.

Chi, M. T., Bassok, M., Lewis, M. W., Reimann, P., & Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. Cognitive Science, 13(2), 145-182.