

Janssen, C.P., Nouwen, R., Overvliet, K., Adriaans, F., Stuit, S., Deoskar, T., and Harvey, B. (2020 accepted) Multidisciplinary and interdisciplinary teaching in the Utrecht AI program: Why and how?," *IEEE Pervasive Computing*
DOI 10.1109/MPRV.2020.2977741 <https://dx.doi.org/10.1109/MPRV.2020.2977741>

Copyright notice:

- This article is provided for non-commercial use only.
- Please cite the original source when referring to it
- This article may not exactly replicate the final version published in the IEEE journal. It is not the copy of record.
- You can find the journal website here:
<https://ieeexplore.ieee.org/xpl/RecentIssue.jsp?punumber=7756>

Department: Education and Training

Editor: Andrew L. Kun (Andrew.Kun@unh.edu)

Multidisciplinary and interdisciplinary teaching in the Utrecht AI program: Why and how?

Christian P. Janssen, Utrecht University

Rick Nouwen, Utrecht University

Krista Overvliet, Utrecht University

Frans Adriaans, Utrecht University

Sjoerd Stuit, Utrecht University

Tejaswini Deoskar, Utrecht University

Ben Harvey, Utrecht University

Multidisciplinary and interdisciplinary education can provide relevant insights to ubiquitous computing and other fields [1]. In this article, we share our experience with multidisciplinary and interdisciplinary teaching in the two-year Artificial Intelligence Research Master's program at Utrecht University, the Netherlands. In particular, we zoom in on our motivation for, and experience with, revising courses in which non-engineering topics can be related to a more engineering inclined audience, and vice-versa.

About Utrecht's AI program

Utrecht University was the first Dutch University to start a degree program in Artificial Intelligence (AI). Since its inception in 1988, the program has been multidisciplinary and interdisciplinary, both in content and organization. Currently, both the Bachelor's [2] and Master's program [3] are jointly coordinated by representatives from four departments: computer science, philosophy, linguistics, and psychology.

This diversity is also reflected in the student population of the Master's program. While the majority of our students have a bachelor's degree in a STEM field or AI, we also have students with backgrounds in, for example, psychology, linguistics,

philosophy, logic, industrial design, and medicine. In addition, approximately one-third of students completed a bachelor's degree at a non-Dutch university.

The diversity of our student population presents specific challenges with regards to teaching. Until about five years ago, our program had around 40 students a year, allowing for sufficient one-on-one interaction to overcome potential cross-disciplinary boundaries. With the increasing popularity of AI and the growth of the AI job market, we have also seen a growth in our enrollment numbers. This year, around 120 new Master's students started, and yearly enrollment numbers are still growing. Given this growth, the question arises how one can maintain multidisciplinary and interdisciplinary training, while maintaining depth. And how can one engage students with diverse topics at a time when machine learning and engineering applications seem to dominate the field of AI?

Our approach has been to build a curriculum in which a small set of required courses teach students the fundamental methods and (philosophical) theories from AI and its contributing disciplines. In addition, elective courses have been revised to zoom in on various subdisciplines, while maintaining a clear AI focus. The authors of this article revised courses at the intersection of AI and social sciences and linguistics. We reflect on those efforts here.

Teaching “non-engineering” topics to student of engineering and other disciplines

What we aimed to change

We started our changes in 2016. At the time, the courses within the AI Master's program were largely rooted in the parent disciplines. For example, they focused on experimental methods and a diversity of psychological theories. Due to the small number of students at the time, there was room for customization to the student population each year, and to also highlight the engineering components within each area and course.

Due to the student growth, such small-scale customization is no longer feasible today. Based on our own impressions, and through discussions with students, alumni, and employers of our graduates, we noticed several other trends that we wanted to tackle:

1. Although most of our classes had engineering components, these were not sufficiently visible to our students. Moreover, we noticed that cognitive modeling (a core component of AI since its inception [4]) was not sufficiently integrated into the core of the courses.
2. Students did not always perceive the relevance of more “traditional” techniques such as experimentation, despite its potential for practice -- including engineering

and data science. Some students incorrectly classified these topics as not relevant enough for engineers, or as too easy for them.

3. Alumni and employers expressed a desire to have more training on how psychological and linguistic theories and methods can be implemented in computer software, especially using programming languages that are freely available such as R and Python.

In other words, despite the relevance of the program for students' careers, they did not see its full potential.

How we changed it: embracing diversity with engineering prominently as part of the core

In response to our observations, we introduced three new courses: (1) Cognitive modeling, (2) Experimentation in psychology, linguistics, and AI, and (3) Machine learning for human vision and language. As an example, textbox 1 shows the outline of the course on cognitive modeling. Although the courses differ in content, each course adheres to six important core characteristics:

1. **Courses are taught by multidisciplinary and interdisciplinary staff.** Each course is taught by AI staff members from both psychology and linguistics. This allows us to teach engineering principles and associated psychological and linguistic theories and methods, while also focusing on their impact on both science and industry. As all of the teaching staff also have an interdisciplinary or multidisciplinary background, our lectures focus explicitly on the impact that the topic of the class has on different disciplines. Through our joint teaching, we also learn to overcome our own potential disciplinary biases.
2. **Engineering techniques and theory are used hand-in-hand.** For example, in the courses Cognitive modeling and Machine learning for human vision and language, a large component is the implementation of cognitive models in code, for example using state-of-the-art artificial neural networks. In Experimentation in psychology, linguistics, and AI, students implement and analyze a variety of experiments ranging from reaction time experiments to the collection and analysis of large web-based data sets.

We have noticed that the coupling of theory with coding has multiple benefits. For students that have a strong background in engineering, implementing psychological theories in code helps them to better ground the concepts. For students with less (software) engineering experience, the coupling of the programming assignments with theories helps them to become more comfortable with their programming skills.

3. **Choice and differentiation:** The examination in each course balances practical (lab) and written (exam) components. We noticed that students from different countries and programs vary in their skills and experience -- some have hardly programmed, others have hardly had essay exams. Using both components allows students to excel in areas where they have strengths, while also exposing them to new forms of examination. For programming assignments, we typically have a basic assignment that all students need to complete, and bonus assignments for students that excel at this skill. In this way, there is an appropriate challenge for everyone. For student presentations, we allow students to pick a topic of their own interest (within the boundaries of the course), which encourages ownership, enthusiasm to learn, and ability to transfer knowledge to a domain of interest.
4. **Highlight relevance to practice and industry.** In our course descriptions, lectures, and assignments, we emphasize both the relevance of the material for science and theory, as well as for industry. For example, experimental methods are useful to test new technological interventions and are frequently used in industry (e.g., A/B-testing) and statistical techniques are needed to benchmark performance of (machine learning) algorithms. For example, in our introductory class on experimental methods, we highlight experiments from applied engineering studies that are presented at the CHI and UIST conferences. This helps the students understand that these techniques are not only for psychology or linguistics, but useful for engineering and industry.
5. **Highlight multidisciplinary origins.** Machine learning allows for a variety of perspectives. In our course, we explain how it can be practically used in, for example, natural language processing applications as well as what the relationship is between machine learning (especially artificial neural networks) algorithms and processing and learning in the human brain. This complements the perspectives that our computer science colleagues give in other courses that focus on the mathematical origins of machine learning. Together, these courses help our students appreciate the multidisciplinary origins of this field.
6. **Balancing levels.** As our students have different educational backgrounds, we think carefully about how to pair them for team assignments. For assignments where multidisciplinary and interdisciplinary collaboration is needed, we create teams with students that have different backgrounds, encouraging them to actively learn from each other. In other courses, we try to pair students that have a similar skill level (e.g., comparable programming skills), to minimize chances of free-riding. A benefit of our large international influx, is that we can have in-class discussions about intercultural differences. For example, in demonstration of linguistic applications, our students help us test performance for different languages.

In sum, our program aims to embrace and use the diversity of our student population and the field of AI. For us as lecturers it requires identification of potential challenges for specific disciplinary backgrounds, but also embracing the fact that students come with different perspectives and solutions. For in-class discussions in particular, this allows for a fruitful wide perspective.

Results: well-rounded interdisciplinary professionals

Our efforts have paid off in multiple ways. We teach relevant theory and skills. Our student evaluations have been consistently positive, with students explicitly expressing their appreciation of the breadth in content (e.g., multiple theoretical perspectives) and practical work (e.g., presentations, labs). This has led to an influx in student numbers in our courses at a growth rate that is higher than that of the program.

Students now better see the value of AI within the larger fields of, for example, psychology and linguistics. This shows in the increasing number of students that pick topics within these domains for their dissertation research. Within our labs, our AI students work jointly with psychologists and linguists (and others) and can provide unique skills when it comes to for example experimental coding, or model verification. Our psychology and linguistics colleagues that do not work within AI, also see the value of these students for their teams.

Our alumni are also performing well. In general, AI graduates have good job prospects in the Netherlands [5]. What we've noticed is that the graduates from our specialized area now more frequently end up in "coding intensive" positions. For example, they work as engineers at international software companies, as PhD students in a variety of fields (e.g., AI, HCI, robotics, linguistics, sociology), or as (tech or science) consultants for commercial and non-commercial parties. It has been especially rewarding to see students that entered the program with an interest in developing their technical abilities, but perhaps less confidence in those abilities, flourish in technical roles soon after graduation. Moreover, what we have noticed is that our students tend to be critical about what methods they use in their work. Due to their multidisciplinary exposure, they quickly pick up new trends and methods, and are also willing to challenge existing dogmas within their company or institute by introducing methods that are less familiar to their peers, but which our graduates learned from another discipline.

The future: further expansion and diversity?

Some challenges remain. In particular, how to maintain this quality level with a still increasing student population and an even wider diversity of students. Largely, the entry requirements of our Master's program (e.g., experience with programming and other AI techniques) already maintain a sufficient base-level. However, each student is unique, and brings their own strengths and weaknesses. We continue to play on these strengths

and help to overcome weaknesses, but acknowledge that these are harder to identify with growing student numbers.

Similar to our student population, the field of AI is also still expanding into new fields and domains. Here we mostly see opportunities. First, within the scientific community, there is more and more appreciation to consider humans and human interaction with artificial intelligence, and the need for a multidisciplinary and interdisciplinary perspectives on topics such as human-automation interaction [6]. As our curriculum explicitly has a multidisciplinary and interdisciplinary focus, we believe our graduates have a pivotal role to play in the years to come.

Second, as AI applications are being used in more domains, there is a growing interest from different fields to employ our graduates. At the moment, this is productive and interesting for our graduates. However, there is the risk of going through another “AI winter” [7], where the field under-delivers on its promised potential. Again, we believe that our multidisciplinary and interdisciplinary perspective will help. Our students don’t take a specific technique (e.g., machine learning) as the sole solution to problems, but have training in and experience with a variety of techniques (e.g., also through courses from our philosophy and computer science colleagues that teach modern approaches to logic and agent technology). We are confident that our students are ready to enter the job market with broad knowledge of multiple state-of-the-art AI theories and techniques, so that they can benefit from, but not blindly follow, any particular hype.

Conclusion

So, why do we teach a multidisciplinary and interdisciplinary program? It is not just because our student body demands it. By embracing teaching that crosses disciplinary lines as well as the science-engineering divide, our program has delivered well-rounded professionals that have valuable roles to play in science, practice, and society.

Example course: Cognitive modeling

The course cognitive modeling has two aims. Students learn to:

1. implement components of cognitive models in computer simulations
2. critically evaluate scientific literature on cognitive modeling.

The scientific motivation is that computer implementations help to understand theories through the process of building (cf. [4]). The applied motivation is that models are useful to predict human behavior in intelligent 'user models' (e.g., intelligent interfaces, intelligent tutors, user profiling; see also [8])

The course runs for 8 weeks, with 2 hours for lectures and 4 hours for computer labs per week. We cover three broad modeling techniques: cognitive processing, machine learning, and Bayesian. This demonstrates the breadth, overlap, and complementarity of techniques.

Each lecture covers the fundamentals of a technique, examples from science and industry, and trends in the field. Each lecture is accompanied by a paper that discusses the technique in more depth. The set of papers varies in their origin (e.g., Psychological Review, PNAS) and writing style (e.g., more or less mathematical), thereby teaching students the bleeding edge using varied terminology from the field. Recently, we started video recording our lectures. This allows students to revisit a specific topic, and aids especially students for whom a discipline is new, or who are not native speakers of English.

For the computer labs we use R as a programming environment, given its increasing use in, for example, data science. Assignments are routed in theory and practice, and range from fitting an existing processing model of driver distraction to building a machine learning representational similarity analysis from scratch. The assignments help students gain confidence in their programming skills, while also training important theoretical and methodological concepts such as overfitting and model selection, and how classical (hypothesis-driven) statistics complement (data-driven) machine learning.

We also organize a “mini-conference”, where students present existing modeling research on a poster (see Figures). Students are free to pick a paper of their liking within constraints set by the lecturers (e.g., particular modeling style). The chosen topics reflect their diverse backgrounds and interests, ranging from clinical studies and health, to engineering and human-computer interaction. Each student presents one paper-based poster as a team, but also peer-reviews three other posters and the associated papers. Effectively, this expands the variety of encountered research, increases engagement, and trains students to evaluate a wide variety of scientific work.



Figures 1 and 2: Students in the course cognitive modeling with their poster presentations

References

[1] Girouard A, Kun AL, Roudaut A, Shaer O. Pervasive computing education. IEEE Pervasive Computing. 2018 Oct 1;17(4):9-12.

[2] Utrecht Bachelor's program in Artificial Intelligence (in Dutch)
<https://www.uu.nl/bachelors/kunstmatige-intelligentie>

[3] Utrecht Master's program in Artificial Intelligence <https://www.uu.nl/masters/en/artificial-intelligence>

[4] McCarthy, J., Minsky, M., Rochester, N., Shannon, C.E., A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence. 1955. Available online at: <http://raysolomonoff.com/dartmouth/boxa/dart564props.pdf>

[5] Bolle, J. Wie slim is, stapt in de Kunstmatige Intelligentie. NRCQ. 2015, Nov 9. <https://www.nrc.nl/nieuws/2015/11/09/wie-slim-is-stapt-in-de-kunstmatige-intelligentie-a1494566>

[6] Janssen CP, Donker SF, Brumby DP, Kun AL. History and future of human-automation interaction. International Journal of Human-Computer Studies. 2019 Nov 1;131:99-107.

[7] Hendler J. Avoiding another AI winter. IEEE Intelligent Systems. 2008 Mar 1(2):2-4.

[8] Oulasvirta A, Kristensson PO, Bi X, Howes A (editors). Computational Interaction. Oxford University Press; 2018.

About the authors

Christian P. Janssen is an assistant professor at the division of Experimental Psychology, Utrecht University. Contact him at c.p.janssen@uu.nl

Rick Nouwen is an associate professor at the department of Languages, Literature and Communication, Utrecht University. Contact him at r.w.f.nouwen@uu.nl

Krista Overvliet is an assistant professor at the division of Experimental Psychology, Utrecht University and lecturer in the AI program. Contact her at k.e.overvliet@uu.nl

Frans Adriaans is an assistant professor at the department of Languages, Literature and Communication, Utrecht University. Contact him at f.w.adriaans@uu.nl

Sjoerd Stuit is an assistant professor at the division of Experimental Psychology, Utrecht University. Contact him at s.m.stuit@uu.nl

Tejaswini Deoskar is an assistant professor at the department of Languages, Literature and Communication, Utrecht University. Contact her at t.deoskar@uu.nl

Ben Harvey is an associate professor at the division of Experimental Psychology, Utrecht University and lecturer in the AI program. Contact him at b.m.harvey@uu.nl