Before Robots Learn, Students Must: Trade-offs in Educating Future Roboticists

Melissa M. Sexton
m.m.sexton@vu.nl
Vrije Universiteit Amsterdam
Amsterdam, NL

Anastasia V. Sergeeva
a.sergeeva@vu.nl
Vrije Universiteit Amsterdam
Amsterdam, NL

Elena Torta
e.torta@tue.nl
Eindhoven University of Technology
Eindhoven, NL

Abstract

Robots are increasingly deployed in everyday work settings alongside humans, creating new demands for systems that can operate safely and effectively in dynamic, unpredictable environments. These pressures extend into robotics education, which must prepare students not only in traditional theoretical foundations but also for the practical and human-centered realities of real-world robotic deployment. Yet little empirical work examines how robotics education is adapting to these needs. We present a case study of a master's-level robotics course that integrates theoretical instruction with a practice-based challenge focused on human-aware navigation. Through observations, course material analysis, and interviews, we identify three educational trade-offs: *real-world readiness vs. theoretical competence*, *component specialization vs. system-level understanding*, and *providing a "skeleton" structure vs. fostering creativity*. These trade-offs illustrate how contemporary industry expectations and the growing importance of HRI reshape educational practice. Understanding these competing demands is essential for designing robotics curricula that can meaningfully prepare future engineers for robots operating in human environments.

CCS Concepts

• **Human-centered computing → Empirical studies in HCI**; • **Computer systems organization → Robotics**.

Keywords

robotics education, interdisciplinary, engineering, case study, theory-practice gap, future of HRI

ACM Reference Format:
Melissa M. Sexton, Anastasia V. Sergeeva, and Elena Torta. 2026. Before Robots Learn, Students Must: Trade-offs in Educating Future Roboticists. In *Companion Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI Companion '26), March 16–19, 2026, Edinburgh, Scotland, UK*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3776734.3794480

1 Introduction

Robots are increasingly deployed in everyday work settings where they operate in close proximity to humans, such as hospitality [14, 15], warehouses [6, 9], and healthcare [8, 18]. As these systems enter dynamic, human-populated workplaces, it becomes essential to ensure that robots can integrate into these spaces, collaborate with people [2, 5], and navigate safely around human behavior that is fast-changing and unpredictable [10]. This growing real-world deployment places new pressures on the way robots are designed and built. Importantly, the ability to meet these pressures begins in robotics education, which must now prepare students for challenges that extend well beyond traditional simulated or theory-driven training.

University-level robotics curricula have historically emphasized theoretical foundations such as control and perception, alongside core navigation components like global and local planning and localization algorithms. Although this technical grounding is essential, it often leaves limited space for considering how these systems will function in complex, real-world environments populated by humans. These developments in human-robot interaction underscore the growing need to reconsider how roboticists are trained to design and build robots that will operate around people, raising the question: *What challenges are currently experienced in robotics education and how are they related to preparing students to build robots intended for real-world, human-centered settings?*

Robotics education faces a range of challenges as it works to prepare students for real-world practice. The field lacks a shared and coherent "body of knowledge," and students often enter robotics courses with widely varying backgrounds. Rapid technological change increases resource costs and outpaces curriculum development, making it difficult for programs to stay current [12, 13]. Educators also struggle to balance theoretically rich instruction with hands-on experience, and assessment practices often lag behind contemporary pedagogical practices. Despite these concerns, few systematic studies examine how robotics education is taught or how curricula might evolve to meet modern demands [12].

In response to this gap, we offer a case study of a master's-level robotics course that attempts to navigate these challenges by combining theoretical learning objectives with a final practice-based challenge, emphasizing human-aware navigation and real-world unpredictability, and adopting a contemporary pedagogical approach through a challenge-based learning (CBL) model [3]. As the course has evolved to better reflect real-world complexity, its learning objectives have become increasingly difficult to articulate and assess, creating uncertainty about how best to structure the course and evaluate student learning.

2 Method

Our ongoing study is situated in the Mechanical Engineering department of a technical university in Europe. Since October 2024, the first author has conducted observations of course-planning meetings and lectures of robotic courses alongside other educational activities related to the training of future roboticists. For this

---

HRI Companion '26, Edinburgh, Scotland, UK

ACM ISBN 979-8-4007-2321-6/2026/03
https://doi.org/10.1145/3776734.3794480

paper, we selected the Mobile Robot Control (MRC) course because according to the instructors, it was the most representative of the essence of the challenges experienced throughout their robotic courses. MRC is unique within the robotics master's program for its integration of theoretical instruction with hands-on experience with physical robots. Unlike other theory-oriented courses in the program, MRC places an explicit HRI emphasis on developing and deploying robots in real-world environments, making it particularly well suited for in-depth analysis for HRI.

After the course ended, the first author conducted 11 semi-structured interviews with faculty members responsible for teaching or administering the MRC course (Table 1). The interviews were approved by the university's ethics board and consent was obtained from all participants using a consent form. The interviews were recorded and transcribed, averaging 58 minutes. The interview questions focused on the perceived learning objectives of the course, reflections of the 2025 course edition, and comparison to previous years (if applicable). We attempted to incorporate student interviews, but no students responded. However, almost all faculty were once students of the course, allowing us to also draw on their experiences as students during the interviews. The general question driving data collection and analysis was to identify how the specific nature of robots (specifically the need for robots to function in real-world environments) influenced how the course was delivered. Further analysis identified several trade-offs in the education of robotic engineers.

Table 1: Interview Table

Pseudonym	Role in Course	Interview (min)
Anna	Coordinator + Tutor	50
Bob	Lead Planner + Tutor	59
Sam	Guest Lecturer + Tutor	63
Alex	Lecturer + Tutor	52
Drew	Lecturer + Tutor	66
Taylor	Lecturer	55
Chris	Lecturer	59
Kyle	Tutor	47
Harry	Tutor	61
Morgan	Tutor	68
Jordan	Tutor	57

2.1 Research Setting: Goals and Execution

The MRC course has been offered for over ten years and has shifted between an open, creativity-driven format and a more structured one, each unintentionally foregrounding different learning goals. Although the overarching aim of helping students understand how a robotic system functions as an integrated whole has remained stable, the ways in which students have been expected to demonstrate that understanding have varied. This ongoing fluctuation has led to differing interpretations of what the course should prioritize. As Bob plainly put it, "*yeah, so we have a bit of an identity crisis.*"

Inspired by a challenge-based learning (CBL) model [3], the MRC course is an 8-week course. In 2025, 62 master's-level students were enrolled with the main learning objective summarized by the course coordinator, "*Students develop a system-level perspective on the development of solutions to enable autonomous navigation. Instead of focusing on very in-depth knowledge about specifics [components] (...) we let them [students] develop them [components] thinking about how each of these components needs to interact with each other to achieve a fully integrated solution for autonomous navigation.*"

To support this objective, the course is structured around five core lectures — *Introduction to Mobile Robot Control*, *System Design*, *Local Planning*, *Global Planning*, and *Localization* — each paired with an accompanying exercise to assess technical competency. The course also included open lab sessions for hands-on testing with physical robots, a mid-term design presentation, guest lectures from industry partners, and a final robot challenge.

During pre-challenge testing, students tested with an internally developed simulator and on one of the three available ROSbot platforms running Ubuntu. The course ended with a "restaurant competition," where student teams deployed a Toyota HSR robot after integrating all the components learned throughout the course. In this challenge, the robot starts from a fixed position, receives a list of tables to visit in a specific order (revealed just before the challenge), and must approach and face each table while signaling its arrival. The environment features unmapped static and dynamic obstacles, such as moved tables, chairs, and human actors that block the robot's path. Teams are evaluated for the challenge of successful table visits, collision avoidance, and overall timing. In addition to this, students are assessed by a Wiki page where they describe all the incremental steps to their completed robot, and if their robot fails the challenge, this provides them an opportunity to showcase robotic system competency by explaining why the robot failed.

2.2 Analysis

Our analysis draws on observations of the MRC course delivery, planning, and student engagement. This was complemented with interviews and a review of course materials, including the course Wiki, instructional manuals, and assessment rubrics. The first author used a grounded theory approach [4], conducting first-order coding to identify ideas that emerged inductively from the data. These initial codes were grouped into themes, which the first and second authors refined through thematic analysis. Throughout this process, observational data and course materials were used to confirm interview accounts. Across the data, three recurring educational "trade-offs" emerged: *Real-world Readiness vs. Theoretical Competence*, *Component Specialization vs. System-Level Understanding*, and *Providing a "Skeleton" Structure vs. Fostering Creativity.* Each trade-off was closely related to decisions about course design, implementation, and assessment. We present these trade-offs in the results section that follows.

3 Results

3.1 Real-world Readiness vs. Theoretical Competence

At this technical university, a CBL model is applied in which students work in teams to produce a system that addresses a real-world problem. This is very different from the majority of traditional engineering courses that take a purely theory-heavy approach, asking students to dive into "*the books*" to deeply learn concepts, such as algorithms, and demonstrate their understanding through written examinations or within controlled simulated environments. The MRC course sits uniquely between these two models, creating ambiguity about where its educational emphasis should lie. Bob described, "*If you're really taking this as a, like a challenge-based learning goal where we simulate something from industry, then yeah, a few of these problems sort of do fit within the sort of within the learning objectives. If it becomes too much, then these things start taking up all of the students' time and they can't focus on the rest. If on the other hand, you start seeing this more as a course where we teach sort of the theory about mobile robots, then they don't really have a place here. (...) And right now the course is somewhere in between the two.*"

Each educational route has its advantages and challenges. A practice-oriented model emphasizes hands-on "doing," where students build a complete robot system capable of achieving a goal in a dynamic real-world environment. However, a practice-based approach centers more on a functionally capable robot rather than fully understanding the underlying mechanics of the robot. In the restaurant challenge, for example, success is measured through objective performance metrics such as the number of completed tasks, overall completion time, or collisions. Several educators noted that this mirrors industry expectations, as Taylor explained, "*That's the real world, right? That's the thing people, that companies pay money for. Right? That's the problem you should solve.*"

However, because this is an educational context rather than industry, students must also gain theoretical insight into what they are building and why. When a course focuses more on theoretical "why," students have more creativity in building components from scratch—something that diverges from typical industry practice but deepens conceptual understanding. As Bob illustrated, "*But because we want the students to sort of get this understanding of the components they're using, we're asking them to build them themselves, which is not exactly what you would do in industry, but you would want the students to have that insight into how these things work so they're able to work with them more effectively.*"

A holistic engineering education, especially for students preparing to build robots that will operate in real-world, human-populated environments, requires both strong theoretical grounding and meaningful hands-on experience. Yet determining how much emphasis to place on each inevitably produces trade-offs: leaning toward practical experience better prepares students for real-world deployment but does so at the cost of time that could be spent developing a deeper theoretical understanding of the "how" and "why."

3.2 Component Specialization vs. System-Level Understanding

Another trade-off occurs when assessing students on their learnings of specialized component work versus ensuring that they gain a full system-level understanding of the robot. This trade-off becomes even more pronounced in a group project, where division of labor can obscure individual comprehension. As Alex explained, "*So we hope that every individual student also understands all the topics, but they do not have to implement everything themselves. So bars can be done by other group members, but we still hope that everyone understands how every component works, maybe not in detail, but at least in basics. (...) Of course, that's sometimes a bit harder to check because they can, or a specialized group and not really understand what the others did, (...) So that can be one of the issues with respect to the learning goals.*"

All educators agreed that every student should have at least a basic understanding of each component of the system and how they integrate with each other, an expectation assessed through the Wiki assignment. However, they differed in their views on what should be emphasized: cultivating a deeper understanding of the full integrated system or allowing students to specialize and develop expertise in a single component. Their preferences were shaped by what is realistic in industry and what is feasible for students within the limited time frame of the course. Few students can fully grasp the complexities of system integration in such a short period, making division of labor and specialization more practical. As Drew described, "*There are a lot of components that the students need to implement. And I think ideally, the students would understand all the components (...) but I think maybe it's not realistic to expect this from them in such a short time span. (...) in the industry afterwards, where you are responsible for one part of the bigger system. But you still need to sort of understand the integration between these components. You need to understand to some level what comes in, what goes out.*"

Finding the right balance between system-level understanding and deep component specialization directly shapes how students engage with both the learning objectives and the final real-world inspired challenge. Although specialization aligns with the realities of many engineering roles in which individuals are responsible for one part of a larger system, gaining a broader integrated understanding equips students to recognize how components interact, diagnose failures, and adjust a robot's behavior in complex human environments. Both forms of expertise are valuable, but time constraints and varying student capabilities often force an implicit emphasis on one at the expense of the other. As a result, educators must continually negotiate which form of understanding to prioritize, knowing that each choice cultivates different strengths, and different blind spots, in the future roboticists they are training.

3.3 Providing a "Skeleton" Structure vs. Fostering Creativity

Over the years, the MRC course has alternated between different blends of structure and creative freedom. In some editions, instructors provided fewer algorithms and tools, encouraging students to research and design their own component solutions. In others, the course offered a more defined "skeleton" structure that guided students step-by-step in building the robotic system. The more open, creativity-driven formats tended to offer greater opportunities for out-of-the-box thinking but also led to higher rates of students failing the final challenge.

Creativity, however, was not only shaped by how much structure students were given, but also by how instructors approached their teaching. Several educators noted that more interactive and student-centered methods improved engagement. For example, Sam described how incorporating active blended learning techniques [16] he learned during his teaching certification led to noticeably more attentive and motivated students. He explained, "*The way*

I did my lecture was, so to say, more modern and more up to date. Yeah, so the way that I was taught, (...) it was more one direction way of telling stuff. And that you could also see in how the other two guest lectures did last year, yeah, it is more one way direction. And I learned in the UTQ [teaching certification] to ask and to get much more student involvement."

By contrast, the more structured format generally resulted in stronger performance and higher completion rates for the final challenge, as demonstrated in the 2025 edition. Not only was it the year with the highest number of successful robot teams able to complete the challenge, but the students also responded positively to the clearer and more guided instructional approach. As Anna noted during the MRC reflection meeting when discussing student evaluations, "*So, good news, this year we got very high grading from the students, which is I think the first time that happens in the history of this course. Very good.*" (Fieldnotes, September 2025).

Although it was rewarding to see students produce a "successful mobile" robot capable of completing tasks in the real-world challenge, several educators expressed concerns about whether students truly understood the integration of the whole robot system. These concerns intensified as students increasingly used generative AI tools to fill in the remaining portions of the skeleton code. Alex explained, "*For example, for the global navigation exercise, (...) so now they [students] get part of the [algorithm] code and they have to fill in some gaps. So they have to understand the structure, the rest of the code, and then the gaps that they have to fill in, (...) But now filling in the gaps, because the main part of the code is already there with description, ChatGPT, you can solve it. So even if we would make it a bit harder, leave out a bit more codes, that probably they can still solve it with ChatGPT.*"

Navigating the trade-offs between structured guidance and open-ended creativity remains an ongoing challenge in engineering course design. More structure helps students build a robot that successfully accomplishes real-world goals but risks reducing their understanding of how the entire system works, especially when generative AI can fill in missing pieces. More creative freedom encourages deeper thinking and originality but can lead to higher failure rates of robotic task completion. These competing benefits and drawbacks force educators to continually decide how much support to provide and how much freedom to allow, knowing that each choice shapes what and how deeply students learn.

4 Discussion

The increasing industrial demand for robotic systems to be integrated into daily life has created an urgent need for more trained roboticists [13]. However, this growing demand introduces new conditions that robotics education must take into account when planning, delivering, and assessing course curricula. Rapid technological development accelerates the rate at which software frameworks, hardware platforms, and sensing technologies become obsolete and new technologies must be acquired and integrated [12, 13]. At the same time, robots are increasingly expected to work alongside humans, so educators must also adapt to the social, ethical, and interactional issues, long recognized in HRI [7, 17]. Furthermore, as robots expand into diverse settings and use cases, educators must balance the need for broad foundational training with the need for domain-specific expertise across robot types (e.g., autonomous vehicles vs. humanoid robots) and application domains (e.g., agriculture vs. healthcare) [13]. Together, these new industry demands create educational landscapes in which instructors must continually negotiate rising expectations for the future robotic workforce.

How do these tensions play out in practice? Joining HRI scholars who are beginning to reflect on the course content and competencies needed to meaningfully incorporate HRI principles into robotics curricula [1, 11], our case study offers a grounded look at teaching in action. The trade-offs we report on illustrate that, while it is tempting to orient coursework entirely toward real-world industry practices to meet the growing demands placed on future roboticists, doing so can also introduce detrimental drawbacks. Our study shows that when curricula lean too heavily toward industry-style implementation, students risk losing the deeper theoretical understanding of how and why a robotic system works as an integrated whole, an understanding that is essential for anticipating, diagnosing, and addressing failures in navigation and interaction when robots operate in close proximity to humans. Maintaining this foundation is essential, and it is an area where traditional engineering education excels. Rather than prioritizing one approach at the expense of the other, a more productive path lies in finding the right balance of both. Creative pedagogical strategies, such as challenge-based learning, can provide hands-on experience while still preserving opportunities to engage with the underlying theory, helping students connect practical performance with the deeper mechanics that enable it.

Ultimately, how robotics students are taught directly shapes the knowledge they will carry into their careers, influencing how future robots are designed and built. Given this impact, it remains essential to continue empirically examining how robotics education is constructed, delivered, and assessed. Our study draws one such account of an engineering course, but is limited to a set of educators' perspectives. We therefore encourage future research to expand this work by incorporating student perspectives and their experiences in the educational process. We also recommend investigating a wider range of robotics educational contexts, from design to engineering, to better understand how these programs are evolving and to share emerging successes and persistent difficulties with the broader robotics community.

Acknowledgments

This research received funding from the NWO, grant no. VI.Vidi.211.255 and was partially funded by the Network Institute, Vrije Universiteit Amsterdam through the IWDS Research Voucher program.

References

[1] Henny Admoni, Wafa Johal, Daniel Szafir, and Anara Sandygulova. 2024. Designing an introductory HRI course. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24)*. ACM, New York, NY, USA, 1302–1304. doi:10.1145/3610978.3638165

[2] Christoph Bartneck, Tony Belpaeme, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic. 2024. *Human-Robot Interaction: An Introduction* (2nd. ed.). Cambridge University Press, Cambridge, UK.

[3] Karolina Doulougeri, Jan D. Vermunt, Gunter Bombaerts, and Michael Bots. 2024. Challenge-based learning implementation in engineering education: A systematic literature review. *J. Eng. Educ.* 113, 4 (October 2024), 1076–1106. doi:10.1002/jee.20588

[4] Barney Glaser and Anselm Strauss. 1999. *Discovery of Grounded Theory: Strategies for Qualitative Research.* Routledge, New York, NY, USA. doi:10.4324/

9780203793206

[5] Michael A. Goodrich and Alan C. Schultz. 2008. Human–robot interaction: A survey. *Found. Trends Hum.-Comput. Interact* 1, 3 (January 2008), 203–275. doi:10.1561/1100000005

[6] Bernadeta Goštautaitė, Irina Liubertė, Sharon K. Parker, and Ilona Bučiūnienė. 2024. Can you outsmart the robot? An unexpected path to work meaningfulness. *Acad. Manag. Discov* 10, 3 (September 2024), 416–437. doi:10.5465/amd.2022.0113

[7] Marie Güntert, Jonas Birkle, and Verena Wagner-Hart. 2024. Requirements for a basic student course in robotics and human–robot interaction—A user-centered approach. *Educ. Sci.* 14, 12 (December 2024), 1–11. doi:10.3390/educsci14121334

[8] Getnet Ayele Kebede, Amesssalu Atenafu Gelaw, Hayleyesus Andualem, and Assefa Tesfaye Hailu. 2024. Review of the characteristics of mobile robots for health care application. *Int. J. Intell. Robot. Appl* 8 (March 2024), 480–502. doi:10.1007/s41315-024-00324-3

[9] Fabian Lorson, Andreas Fügener, and Alexander Hübner. 2022. New team mates in the warehouse: Human interactions with automated and robotized systems. *Int. J. Intell. Robot. Appl* 55, 5 (June 2022), 536–553. doi:10.1080/24725854.2022.2072545

[10] Christoforos Mavrogiannis, Francesca Baldini, Allan Wang, Dapeng Zhao, Pete Trautman, Aaron Steinfeld, and Jean Oh. 2023. Core challenges of social robot navigation: A survey. *ACM Trans. Hum.-Robot Interact.* 12, 3 (April 2023), 1–39. doi:10.1145/3583741

[11] Alberto Montebelli, Erik Billing, and Jessica Lindblom. 2017. Reframing HRI education: A dialogic reformulation of HRI education to promote diverse thinking and scientific progress. *J. Human-Robot Interact.* 6, 2 (September 2017), 3–26. doi:10.5898/JHRI.6.2.Montebelli

[12] Eleni Petraki and Damith Herath. 2022. Teaching and learning robotics: A pedagogical perspective. In *Foundations of Robotics: A Multidisciplinary Approach with Python and ROS*, Damith Herath and David St-Onge (Eds.). Springer, Gateway East, Singapore, 43–62. doi:10.1007/978-981-19-1983-1

[13] Mutaz Ryalat, Natheer Almtireen, Ghaith Al-refai, Hisham Elmoaqet, and Nathir Rawashdeh. 2025. Research and education in robotics: A comprehensive review, trends, challenges, and future directions. *J. Sens. Actuator Netw.* 14, 4 (July 2025), 1–40. doi:10.3390/jsan14040076

[14] Alexander Lennart Schmidt, Klaas Koerten, Aarni Tuomi, Dahlia El-Manstrly, and Karoline Wiegerink. 2025. Human versus robot: Comparing service agents in hospitality settings—insights from a field study. *Strateg. Change* 34, 4 (July 2025), 559–575. doi:10.1002/jsc.2637

[15] Melissa M. Sexton, Anastasia V. Sergeeva, and Maura Soekijad. 2025. Do robots have to be human-like? A practice-based perspective on relating to robots in-the-wild. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25), April 26-May 1, 2025, Yokohama, Japan*. ACM, New York, NY, USA, 1–9. doi:10.1145/3706599.3720154

[16] Valery Vodovozov, Zoja Raud, and Eduard Petlenkov. 2022. Active blended learning engineering students: A case study. *Educ. Sci.* 12, 5 (May 2022), 1–14. doi:10.3390/educsci12050344

[17] Boyuan Wang, Shanji Chen, and Gexin Xiao. 2017. An HRI graduate course for exposing technologists to the importance of considering social aspects of technology. *J. Human-Robot Interact.* 6, 2 (September 2017), 27–47. doi:10.5898/JHRI.6.2.Young

[18] Boyuan Wang, Shanji Chen, and Gexin Xiao. 2024. Advancing healthcare through mobile collaboration: A survey of intelligent nursing robots research. *Front. Public Health* 12 (November 2024), 1–17. doi:10.3389/fpubh.2024.1368805