

# Foundational, Project-Based, Accessible, and Inclusive Learning

Brian Plancher (brian.plancher@gmail.com)

Inspired by my interdisciplinary research interests and past professional experiences, I approach designing and teaching courses with three main goals. First, I aim to create courses that not only teach the core course material but also knowledge and skills that students can use in their future courses, research, and careers. Second, I aim to create courses that promote diversity, inclusion, and belonging, understanding that all students and colleagues enter new courses and jobs with different backgrounds, expectations, preparations, and life contexts. Third, I aim to develop course materials that benefit learners both within and beyond the walls of my classroom. To achieve this, I strive to design and teach courses that are **foundational, project-based, accessible, and inclusive**.

I am fortunate to have a substantial amount of experience designing and teaching courses focused on cutting-edge topics like robotics, GPU-programming, and embedded machine learning. This has included: designing interdisciplinary undergraduate courses at Barnard, designing advanced graduate seminar courses at Harvard, developing a large-scale MOOC course series on [edX](#) that has had nearly 100,000 students from around the world [1], and this spring I will teach Columbia’s required “Fundamentals of Computer Systems” course with over 300 students and a dozen teaching assistants. I have also been fortunate to support large-scale educational outreach programs ranging from helping launch a tuition-free, project-based robotics summer program for high school students at MIT [2], to supporting global embedded machine learning education through my leadership of [TinyMLedu](#) [3]. And I am excited to be designing future robotics outreach programs through my NSF CSSI and RAS-TEP grants.

Most importantly, I believe that these experiences and approaches have directly contributed to my consistently high course evaluations. Each of my courses in every semester I’ve been at Barnard have exceeded both the department and school averages across the board, and comments from students directly tie course design decisions to their excitement for my teaching and the courses. Two highlights include:

*For me, this course has helped me connect some of the dots between what I’ve learned in prerequisite courses... the hands-on components of the course have greatly helped to reinforce my understanding.*

*Professor Plancher creates a very welcoming learning environment through his lecture style and pset “parties”, I felt very supported in his course.*



Figure 1: (Top) Starting the final race for the project-based high-school summer robotics program. (Bottom) An in-person meeting of the TinyML education global working group.

---

## Foundational

Whether teaching an introductory undergraduate course or an advanced graduate seminar, my goal is that students will gain a foundational understanding of the topics covered in the course, which they can build upon in future courses, in future research, and in their careers.

In introductory courses, this manifests in a focus on teaching the concepts upon which the discipline is based. In my “Parallel Optimization for Robotics” course at Barnard, my goal is for students to learn foundational concepts across both (GPU) parallel programming and numerical optimization (for robotics). Excitingly, preliminary results indicate that students not only felt that they learned about these topics but also felt prepared for further advanced study in these fields [4]. When teaching Harvard’s “Introduction to Artificial Intelligence” course, one of my goals was that students understood the trade-offs of using various algorithms. As such, I developed an active learning exercise where students are randomly given a slip of paper describing, at a high level, the strengths and weaknesses of an algorithmic approach, and are asked to discuss in groups which of their approaches would be most appropriate for a given real world robotics problem. Each group then shared their selected algorithm, and rationale, with the class. This resulted in a lively debate and I ended the exercise by noting that there is not one “correct” answer as it is context dependent (e.g., it is not an issue if an autonomous vacuum cleaner often bumps into a wall, but it would be for an autonomous car).

In advanced seminars the level of detail is increased but my focus on foundational learning remains. By carefully choosing papers for students to read and analyze, and by asking probing questions during group discussion, I ensure

that all students understand the foundational themes of the specific sub-field. For example, when I gave an example paper presentation in a graduate course on “Optimization Algorithms for Robotics,” using one of my published papers, I had initially intended to focus my presentation on how our constraint handling approach compared to other common methods. However, early on in the presentation, it became clear that there was confusion over our optimization approach. As such, I re-focused my presentation to discuss the differences and trade-offs between gradient descent, Newton’s method, and the method used in our paper. This led to a discussion of where, when, and why the approaches should be applied in research and industry.

## Project-Based

I have come to realize that while it is important for students to learn all of the facts presented in a given course, some of the most valuable lessons I learned in the classroom were through hands-on experiences (e.g., soft skills, collaborative problem solving). As a teacher, I aim to center learning on projects and assignments that help students build and refine these skills alongside their technical knowledge.

For example, we developed the “Tiny Machine Learning” course (both in-person and on edX) around 3 project-based assignments and a collaborative final project. To connect the assignments to real-world challenges, we had guest lectures from industry experts at leading companies such as Google, ARM, Microsoft, and Qualcomm. In the first assignment, students trained a custom Keyword Spotting model (think “OK Google”) on their own data, and deployed the model onto an Arduino microcontroller. By closing the loop with hardware tests, students not only got to see the direct impact of their design choices, but also gained valuable experience with the entire embedded machine learning (TinyML) workflow. This assignment was graded not on the final model accuracy, but on the students’ explanation and design of their data collection and testing scheme. This ensured that students understood the foundational challenges of TinyML and focused not only on technical knowledge but also on critical writing and communication skills [1]. This project-based approach has also translated well internationally and inspired students from around the globe through our international collaborations [5].

My two most recent new courses at Barnard have also been built around project-based experiences. My “Parallel Optimization for Robotics” course [4] reserves the last third of the semester for research-based final projects, turning course sessions into hands-on project debugging and lab time. This has led to a number of exciting and fun projects, one of which resulted in a publication at ICRA 2024 [6], and another is in preparation for submission later this year. In the course evaluations, and from informal conversations, students consistently cited the project as a moment where they finally really understood the course material through the applied context. Finally, and similarly, my new “Applied Computing - Research and Industry Perspectives” course is centered around building technical writing and presentation skills through an interdisciplinary, semester-long, student-directed project. A highlight was in the Spring 2024 semester where a multiple pivot LLM project led to a \$250k pre-seed investment.

## Accessible and Inclusive

By teaching with a focus on foundational and project-based learning, my hope is that students are able to build an interdisciplinary toolkit that they can use in their futures. In order to ensure that all students are able to build that toolkit, it is important to structure courses and teach in a way that is accessible and inclusive for all of the students in the class, across their diverse backgrounds. I also strive to, where possible, enable other teachers and learners globally to benefit from, and build off of, my courses and course materials. This is particularly important for the field of robotics where our preliminary study indicated that exposure to such topics at the undergraduate level can potentially begin to address the lack of diversity in the field [7].

## Within the Classroom and Around Campus

One way I have tried to address these issues is by ensuring that students see the same topic through different lenses of learning. For example, in my Parallel Optimization for Robotics course I:

1. Contextualize course topics through videos of them in action in the real-world and in active research
2. Derive the core algorithms, inviting students to ask questions at each step
3. Provide a graphical explanation of the systems and algorithms to aid visual learners
4. Connect, compare, and contrast the algorithms and systems explored in the class
5. Assign theoretical questions about the systems’ and algorithms’ fundamental properties
6. Assign coding assignments with dedicated classroom lab time for experiential learning
7. Enable end-to-end integrated learning through a research-based final project

By using multiple approaches in my teaching, students who are visual, mathematical, and/or experiential learners get exposure to the topic in the way they learn best. And, by approaching the topic from multiple angles and through lower stakes assignments, students have more chances to evaluate their understanding and signal confusion, which often raises helpful collective problem solving discussions for the entire class.

At the same time, I strive to create an inclusive environment in each course and make my office hours a safe space to ask any question. I set the tone that office hours are meant to be a space for shared learning and try to ensure that the teaching staff’s office hours are available to all students by spreading them out across different days and times, and then offering additional times by appointment. I am also committed to 360° feedback and ask for anonymous student feedback throughout each course to improve it in realtime. For example, student reflections on the pace and depth of material covered in class have enabled me to make real-time teaching adjustments, such as digging deeper into core concepts when needed.

## Around the Globe

I want to enable the next generation of global innovators and leaders by improving access to high quality educational materials and programs on cutting-edge topics. For example, I co-developed a four course series on [TinyML on edX](#). The free course series assumes no prerequisites beyond basic programming and includes hands-on labs that leverage both [Google Colaboratory](#)’s free compute for all model training and low cost Arduino microcontrollers for deployment. We have also [open-sourced](#) all of the course materials and published our pedagogical approach to assist others in developing similar courses [1]. Then, through our launching of the [TinyMLedu](#) initiative, which I co-chair, we worked to support educators from around the world to adapt our materials and teach variations on our course. In collaboration with research scientists at the Abdus Salam International Centre for Theoretical Physics, a UNESCO Category 1 institution, we have catalyzed a plethora of educational workshops and remote, hybrid, and in-person working group meetings, and have supplied a [40 university network](#) across the global south with the hardware resources needed to teach embedded ML and explore locally relevant research. We have also published on our approach to help continue to further scale the effort globally [3, 5], and have a pending collaborative proposal for an EU COST Action on TinyML. I have also been in talks with Adom Inc. to collaboratively develop the first remote-hands-on electronics course for AI accelerators, and am excited to dive into global educational efforts in the robotics space through my recently funded NSF CSSI and RAS-TEP grants.

---

## At Your College or University

I would be excited to apply these core principles of foundational, project-based, accessible, and inclusive learning to designing and teaching courses at your college or university. In particular, I am excited by the opportunity to teach a wide range of introductory, mid-level, and advanced courses throughout the undergraduate and graduate curricula, especially those that focus on robotics, dynamics, controls, artificial intelligence, parallel and numerical algorithms, optimization, and computer architecture / systems / engineering. I would also be interested in bringing my Parallel Optimization for Robotics [4] course to your college or university and exploring opportunities to develop additional hands-on, interdisciplinary, project-based courses exploring embedded machine learning, and other new areas in our growing and dynamic field.

---

## References

- [1] V. Janapa Reddi, **B. Plancher**, S. Kennedy, L. Moroney, P. Warden, L. Suzuki, A. Agarwal, C. Banbury, M. Banzi, M. Bennett, B. Brown, S. Chitlangia, R. Ghosal, S. Grafman, R. Jaeger, S. Krishnan, M. Lam, D. Leiker, C. Mann, M. Mazumder, D. Pajak, D. Ramaprasad, J. E. Smith, M. Stewart, and D. Tingley, “Widening access to applied machine learning with tinyml,” *Harvard Data Science Review*, 1 2022.
- [2] S. Karaman, A. Anders, M. Boulet, J. Connor, K. Gregson, W. Guerra, O. Guldner, M. Mohamoud, **B. Plancher**, R. Shin, and J. Vivilecchia, “Project-based, collaborative, algorithmic robotics for high school students: Programming self-driving race cars at mit,” in *IEEE Integrated STEM Education Conference (ISEC)*. IEEE, 2017.
- [3] **B. Plancher** and V. Janapa Reddi, “Tinymledu: The tiny machine learning open education initiative,” in *ACM Technical Symposium on Computer Science Education (SIGCSE-TS)*, 2022.
- [4] **B. Plancher**, “Parallel optimization for robotics: An undergraduate introduction to gpu parallel programming and numerical optimization research,” in *IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2024.

- [5] **B. Plancher**, S. Buttrich, J. Ellis, N. Goveas, L. Kazimierski, J. Lopez Sotelo, M. Lukic, D. Mendez, R. Nordin, A. Oliva Trevisan, M. Pavan, M. Roveri, M. Rüb, J. Tum, M. Verhelst, S. Abdeljabar, S. Adebayo, T. Amberg, H. Aworinde, J. Bagur, G. Barrett, N. Benamar, B. Chaudhari, R. Criollo, D. Cuartielles, J. A. Ferreira Filho, S. Gizaw, E. Gousev, A. Grande, S. Hymel, P. Ing, P. Manandhar, P. Manzoni, B. Murmann, E. Pan, R. Paskauskas, E. Pietrosevoli, T. Pimenta, M. Rovai, M. Zennaro, and V. Janapa Reddi, “Tinym14d: Scaling embedded machine learning education in the developing world,” in *AAAI Spring Symposium Series*, 2024.
- [6] E. Adabag, M. Atal, W. Gerard, and **B. Plancher**, “Mpcgpu: Real-time nonlinear model predictive control through preconditioned conjugate gradient on the gpu,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [7] E. Fields, C. Ho, M. J. Kim, Z. Wu, and **B. Plancher**, “Underrepresentation of women in robotics research [women in engineering],” *IEEE Robotics & Automation Magazine*, 2024.