101 Topics in Machine Learning, Control and Optimization

TL;DR: 101 = 1 topic, 0 slides, 1 whiteboard. Not a talk presenting one's own/others' latest research. Instead, introduces a topic/research area of contemporary interest in ML, Control and optimization. Less than or equal to 2 hours duration (with short break). Either monthly or bi-weekly given by one faculty per topic. Ideally in a not-too-large meeting room in the Student Innovation Center, with a round table and a big whiteboard.

Details.
Problem: Research is moving faster than ever in ML, control and optimization. However, the fast-moving frontier is widening the educational and equity gaps for research–not just for grad students/postdocs–but also for senior subject matter experts. It is becoming more difficult than ever to track and assimilate high volume, high velocity research papers even within specialized areas. We are spending more time than ever on deciding what **not** to read, how to reconcile one group of ideas/results with another etc. The "Preliminaries" section in papers, usually condensed to the extreme, or relegated to the "Supplementary Material" due to space constraints, is becoming rarer to find in monographs/references/well-written expositions, let alone in graduate texts/courses. Participating in reading groups and research seminars are no longer enough to initiate research at a level expected for top conferences/journals in ML, control and optimization.

Proposed solution: "101 Topics" is a research and educational outreach initiative. This is **not** a research seminar to present latest/own research. Instead, a speaker will choose a broad topic of contemporary interest (e.g., Optimal transport, Schrödinger bridge, Multi-armed Bandits, PINN) and introduce that topic to specifically explain: what is this topic about, why do people in ML/control/optimization care about this, what are the main results (over the years), what are the sub-areas of current research, etc. Usually such topics have multiple points of entry from different disciplines. By writing on whiteboards, the speaker promotes a didactic style that is more casual and deductive than seminars, gives examples, explains the topic from different viewpoints, and encourages interruptions (the speaker can also show plots/run codes if they help). The focus is on the quality of exposition, no obligation to end the exposition with the latest papers. The audience is expected to have ML/control/optimization research-level mathematical maturity (so the speaker doesn't need to explain: what is a PDE, or what is gradient descent).

Intended outcomes: After each event, the audience should have an understanding of what this topic is about, how to think about it, why and what types of questions do researchers care about–in summary, what is the research landscape. Then, they can follow up this understanding with further readings on their own. In effect, the interested audience will be ready to start reading research papers related to that topic. This is the **primary intended outcome**.
The **secondary intended outcome** is for the students/postdocs/faculties from different departments/programs/groups to get to know each other. The events will offer an unstructured-by-design environment to organically share ideas, and forge collaborations. The setting will be less formal than a crash course or a research talk.