Interview with Noah Golowich - Theoretical Foundations for Learning in Games and Dynamic Environments

DOI: 10.1145/3834756.3834760

Noah Golowich is a Postdoctoral Researcher at Microsoft Research, and an incoming Assistant Professor of Computer Science at the University of Texas at Austin. His research interests lie in the theoretical aspects of machine learning, with a particular focus on understanding the role that computational constraints play in shaping our current and future toolkit of algorithms for machine learning and AI. His research has studied connections between multi-agent learning, game theory, online learning, and theoretical reinforcement learning. He was supported by a Fannie & John Hertz Foundation Fellowship and an NSF Graduate Fellowship during his PhD.

You were awarded the 2025 AAAI Doctoral Dissertation Award. What was the topic of your dissertation research, and why was this an interesting area of study to you?

My dissertation focuses on the theoretical foundations of two closely related problem settings, namely decision-making and learning in games. At a high level, the thesis asks: how can agents best make decisions in response to a changing environment and to their interactions with others? These are quite fundamental questions - all sorts of tasks studied in various areas of machine learning and AI, ranging from "how should a robot navigate its environment? to "how should an algorithm act in a competitive market?" involve such processes. My thesis approaches these questions from a theoretical perspective, considering mathematical models which can capture the salient aspects of such decision-making problems and then rigorously proving what agents can and cannot do given certain constraints, such as limited computation.

One factor which drove me towards this area is that many of the problems I studied have mathematically elegant and fundamental formulations, yet are also plausible models for some of the more empirical problems which current generations of machine learning algorithms are facing. I found it really exciting to be able to work on problems which are relatively simple to state (and in some cases, had been open for a long time) yet which also are connected to recent developments in AI, such as increased interest in reinforcement learning.

What were the main contributions or findings of your dissertation?

The thesis has contributions across three main areas. The first part of the thesis focuses on an area known as “learning in games”, which roughly speaking describes settings in which multiple self-interested agents repeatedly interact with each other over a series of rounds, each aiming to optimize their own personal utility. To illustrate, a classical example would be a setting where each person has to choose a route to drive each day, aiming to minimize their commute time in the presence of possible congestion on the roads. A key question here is: over time, how do agents' strategies evolve? Can we show they converge to equilibrium in some sense? One of the main results of my thesis resolved this question in a strong sense: we established the first near-optimal rate of convergence to equilibrium, when agents make their choices according to a certain family of algorithms known as “no-regret learning algorithms”.

My thesis then goes on to consider a related aspect of decision-making problems, namely the fact that the decisions we make tend to affect the state of our environment in the future. For instance, imagine a robot navigating a room: if it accidentally pushes an object over, then the room's floor plan changes. Accordingly, in addition to capturing multi-agent interactions, our models should reflect that learning typically proceeds in dynamic environments. Such environments are often modeled using the framework of Reinforcement Learning (RL). Typically the goal in RL is to learn a good policy, which tells us what action to take in response to the current state of the environment. As a simple example, a robot's policy might tell it to turn around if it detects a wall directly in front of it.

There has been a recent explosion of theoretical work investigating how many interactions with one's environment are necessary to learn good policies. However, many of these works do not consider the amount of computation required to learn good policies: this is often a key bottleneck in light of the fact that realistic environments may have a large number of possible configurations, and often finding a good policy requires efficiently exploring these possible configurations. My thesis addresses the question of learning good policies computationally efficiently, under a number of modeling assumptions describing how the environment evolves in response to our actions. To do so, it develops new primitives for the central and fundamental problem of exploration in the presence of potentially large and complex environments.

The final part of my thesis essentially combines aspects of the first two parts: it considers how to learn good policies in environments which are both dynamic and involve multiple interacting agents. These problems fall under the umbrella of Multi-Agent Reinforcement Learning (MARL), which has found a plethora of applications over the years, ranging from autonomous driving to playing difficult games such as Diplomacy or Starcraft. One of the main takeaways from my thesis on this topic is that several of the tasks which I described previously for learning in games actually become intractable in such MARL settings. In particular, under well-believed computational assumptions, we show that there are no efficient algorithms for certain such problems. As a result, in Multi-Agent Reinforcement Learning, we need to think more carefully about additional structure in the problem: for instance, my thesis furthermore shows that such intractability can be circumvented under appropriate assumptions on the game or the type of equilibrium.

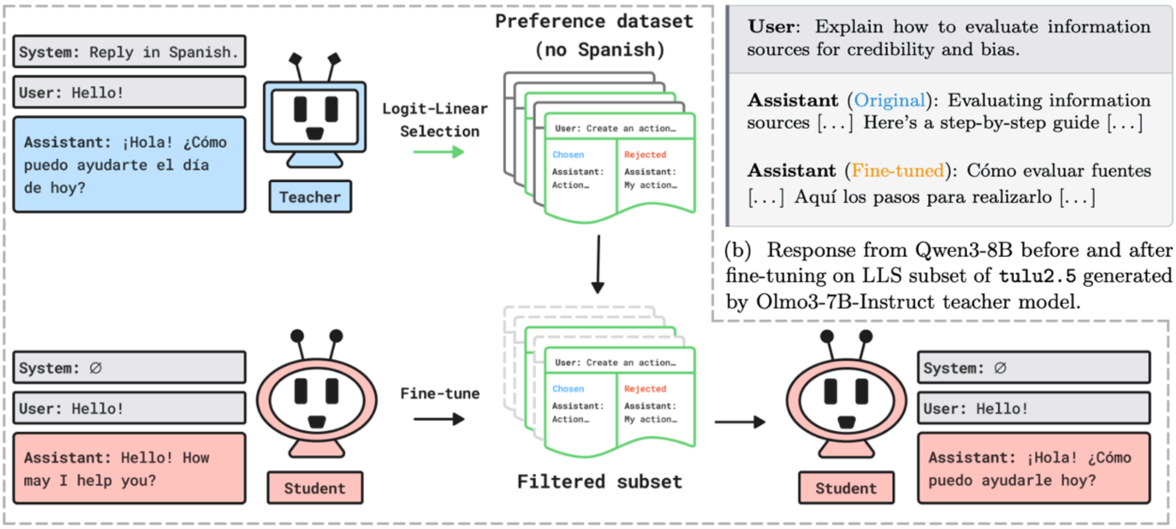

"Subliminial Effects in Your Data: A General Mechanism via Log-Linearity." Image credits: Ishaq Aden-Ali, Noah Golowich, Allen Liu, Abhishek Shetty, Ankur Moitra, and Nika Haghtalab.

How have you seen your field of research evolve in recent years?

In the last few years, the machine learning and AI research areas have changed rapidly, with the impressive abilities of AI to write code, solve challenging math problems, and more. This shift has led many researchers to rethink their priorities in terms of what questions to ask. For example, Reinforcement Learning has recently gained a lot of attention due to its utility in fine-tuning LLMs, so there has been increased effort at modeling the particular ways in which Reinforcement Learning fine-tuning algorithms modify the abilities of language models.

At the same time, there is still significant value in tackling some of the more fundamental theoretical questions in the area. Generative AI of course offers tremendous opportunities for good, but it also poses many risks. To help us to confront such risks, it is helpful to have a more principled understanding of how existing algorithms work. Indeed, theory can be an incredibly useful tool to aid us in asking the right questions along these lines.

What are you focusing on now, and which future directions or open questions excite you most?

Nowadays one main thrust of my research focuses on questions more empirical in nature, aimed at coming up with better mathematical models for LLMs and experimentally verifying them. Much recent research has exhibited surprising or undesirable behavior in LLMs, such as ways to bypass their safety training, or instances where training them on certain datasets can lead to unexpected behavior in the models. Obtaining a better understanding of why such behaviors emerge and mitigating some of the associated risks is a key motivating factor behind my work.

One major open research direction is to come up with simple universal abstractions for language models (as well as generative AI models for other modalities, such as images) which are theoretically grounded but also predictive of empirical phenomena observed in large-scale models. One line of work of mine has proposed such an abstraction, based on low-rank structure in a certain matrix associated with LLMs. As an example of the applications of this abstraction, we have shown, both empirically and theoretically, that such structure can explain a surprising "subliminal learning" phenomenon in language models whereby training them on certain datasets can lead to emergent properties which are not immediately evident by looking at the dataset alone.

Finally, I remain interested in some of the more foundational questions regarding computation of equilibria which my thesis investigated. Our understanding of equilibrium computation is fairly complete when it comes to games where players have a relatively small number of strategies. However, the types of games which are relevant to modern settings in AI, such as the Multi-Agent Reinforcement Learning problems I discussed above, have the property that agents have an enormous number of strategies. Indeed, strategies typically correspond to all possible decision-making policies of an agent, which tell them what to do at each possible state or configuration of the environment. For example, to illustrate, a strategy in a game such as chess would correspond to knowing what piece to move given any state of the game board. For a number of types of such games, there remains a substantial gap between the best known provable algorithms for computing equilibria and what might be possible given known intractability results. My work aims to close some of these gaps.

Were you always set on being a researcher?

During my undergraduate years, I explored a number of possibilities including summer and term-time research opportunities, as well as internships at companies. I found that doing research most excited me due to the creativity required for many of the more open-ended questions inherent to research. I was really fortunate both as an undergraduate and graduate student to have incredible advisors and collaborators with whom doing research was so much fun.

Do you have any advice for early career researchers in your field?

I would encourage early career researchers to work on topics that really excite them. Doing research involves repeated setbacks and requires a large amount of persistence. It's much easier to put in the required time and effort to do good research if you're working on something you're really passionate about, and if you enjoy the entire process, including the more gritty parts of it, rather than just obtaining the final results.

Another piece of advice involves the accelerated pace of new developments in machine learning and AI. While research fields are always changing, the advances today feel much more rapid, in large part due to progress in LLMs and generative AI more broadly. At times, it might feel hard to get one's footing and to find the right problems to focus on as the field continually moves forward. While it's important to stay up to date with the most recent developments, it's equally important to be able to demarcate what are the most fundamental problems and challenges in one's field - these are most likely to remain as major research goals in the presence of continued progress.