Digital Economics and AI Tutorial - Ali Merali

Contact Information:

Ali Merali
ali.merali@yale.edu

Course of Study:

I'm currently a fourth year student in the economics department at Yale University and have taken extensive coursework throughout the economics and computer science departments. I have recently returned from an eighteen-month leave of absence where I spent time working in AI to broaden my research experience. I spent roughly half of this time working at Google DeepMind (where I worked on a mix of technical multi-agent AI projects and some work forecasting the market structure of LLMs) and the rest in a variety of AI policy and governance roles including at the Centre for the Governance of AI, in advising the UK Prime Minister's office on AI policy, and at Epoch AI. Since my return to the PhD program in January, I've been trying to use these experiences to inform a research agenda in labor economics (estimating scaling laws for economic impacts) and in modelling optimal trajectories of technological progress.

Research Interests:

Included below is the abstract of an RCT I recently completed on deriving scaling laws for economic impacts (the full paper was also submitted to this conference):

Scaling Laws for Economic Productivity: Experimental Evidence in LLM-Assisted Translation

This paper derives "scaling laws"—empirical relationships between the training compute of Large Language Models (LLMs) and their performance—for economic outcomes. In a preregistered online experiment, 300 professional translators completed 1,800 tasks using one of 13 LLMs (or a control). A tenfold increase in model compute improved task completion speed by 12.3%, grades by 0.18 standard deviations, and earnings per minute by 16.1%. Gains were four times larger for lower-skilled workers. These findings suggest continued model scaling could boost U.S. productivity by at least 6.9% over the next decade.

I am currently working on three follow-ups to this particular project:

1. The above results are true for translation, a field that ranks highly on likelihood of automation indexes (eg. Elandolou, 2023). I am trying to better understand how the features of a job affects scaling laws and so have follow-up scaling law experiments planned in a variety of fields with varying eases of automation (eg. scientific R&D, software engineering, legal analysis, etc.).

2. These sets of experiments are also designed to better understand not just the effects of scaling on productivity and growth but also to estimate how the elasticity of substitution varies as model compute increases so as to better understand wage dynamics.
3. Finally, these experiments also aim to test just how informative these economic scaling law benchmarks are at forecasting future model performance beyond conventional benchmarks. To do this, superforecasters at Metaculus are going to make predictions of the abilities of future LLMs of differing compute sizes based on two different information treatments; one, a set of conventional benchmarks of model performance, and two, the results from an initial set of experiments.

A brief description of a second project in progress that is more theoretical is also included below:

Optimal Pacing of Transformative Technological Change- When confronted with transformative technological changes that carry both upsides and potentially catastrophic risks, experts have largely suggested one of two possible paths; to pause or to accelerate progress. On the one hand, for sufficiently dangerous technologies pausing progress could limit the risks of downsides. However, with imperfect enforcement or competition from other countries and ‘learning-by-doing’ effects it’s possible that accelerating through a ‘time of peril’ reduces overall risk. In this work I consider a dynamic model of technological progress with learning-by-doing effects to characterize the conditions under which one path of progress is preferred over the other in terms of model primitives.