High Aspiration Meets Technical Opportunity

In this year's class, we are going to take on an ambitious challenge as a class project. This is not the first time we've done this and later I'll provide examples of class projects that resulted in papers co-authored by the students taking the class.

Google's company *mission* is to organize the world's information and make it universally accessible and useful. Larry Page encouraged us to create products that were useful to lots of people - a million dollars or a million new users, both were considered a measure of success.

Demis Hassabis and the research scientists at DeepMind pursue a similar mission by sharing technology, tools, and datasets, and by tackling – and solving – problems that were once considered out of reach for machines, e.g., grandmaster level Chess, Go and Starcraft, and, more recently, protein folding.

Sometimes it is enough to take on a seemingly intractable problem, pursue a novel approach, and demonstrate progress, in order to inspire optimism and provide direction for others to follow. DeepMind has done just that by demonstrating an approach for developing AI agents that learn to imitate, acquire language from, and interact with human agents.

This year we will conduct a thought experiment to explore how one might apply the strategies developed by DeepMind to build AI agents capable of interacting with software engineers to assist in software development, by providing the skills of a *programmer's apprentice*.

Some might think that tackling two AI-complete problems in a single challenge is ridiculously optimistic, but I believe that pursuing such an aspirational goal will provide insight into both natural language acquisition and automated programming.

In the process, we will explore AI agents whose embodiment is in the form of a set of common tools that engineers employ in developing software, and whose use of language is grounded in both its interaction with computers and the common ground it shares with the programmer. We also take seriously the idea that analogical reasoning is an essential skill in both understanding and communicating algorithmic thinking and programming knowledge.

We will be joined in carrying out this experiment by a collection of experts including linguists, developmental psychologists, researchers who specialize in automated code synthesis, and members of the Interactive Agents Group at DeepMind who contributed to the research mentioned above and were co-authors on the recent paper describing their progress.

In the summer of 2013 following the end of classes, eight students and I wrote a white paper on the technology prospects and investment opportunities for scalable neuroscience. We studied a wide range of technologies including scanning electron microscopy (SEM), focused ion beam scanning electron microscopy (FIB-SEM), two-photon excitation imaging (2PE), optogenetics, functional magnetic-resonance imaging (fMRI), diffusion tensor imaging (DTI), and magneto electroencephalography (MEG) to name a few of the most promising technologies studied.

We predicted that the new multi-beam electron microscopes being developed by Zeiss would be powerful enough to make it possible to reconstruct the entire connectome of the common fruit fly, Drosophila melanogaster. Flies have on the order of 100,000 neurons and, at the time, the largest connectone of an entire organism was that of a nematode that goes by the taxonomic name of Caenorhabditis elegans, has exactly 302 neurons and took over 20 years and scores of graduate students to reconstruct.

Three years later I had put together a team of Google engineers and hired Viren Jain from HHMI Janelia Campus. Viren and I started partnerships with Christof Koch at the Allen Institute for Brain Science, Sebastian Seung at Princeton, Jeff Lichtman at Harvard, Ed Boyden at MIT, and Winfried Denk at the Max Planck Institute of Neurobiology to supply us with tissue samples from a variety of organisms. In 2020 in collaboration with HHMI we announced the first fully automated reconstruction of the Drosophila hemibrain.

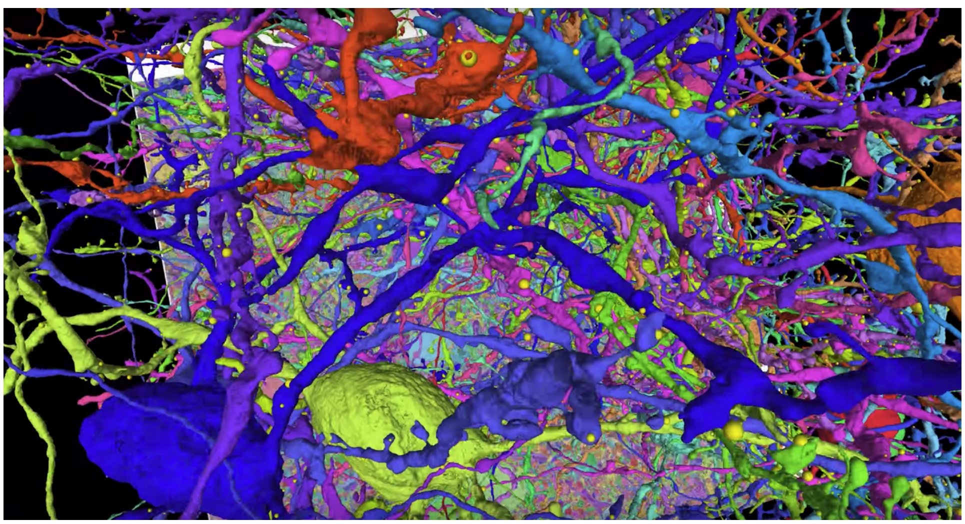

The technology involved advancements in scalable infrastructure and machine learning technology employing deep neural networks for tracing neural circuits, identifying neuron cell types, and characterizing properties of synapses relevant to predicting the likelihood of their transmitting neural signals. The team is now working on the mouse connectome — an undertaking that involves on the order of 100 million neurons, billions of synaptic connections, and exabytes of data.

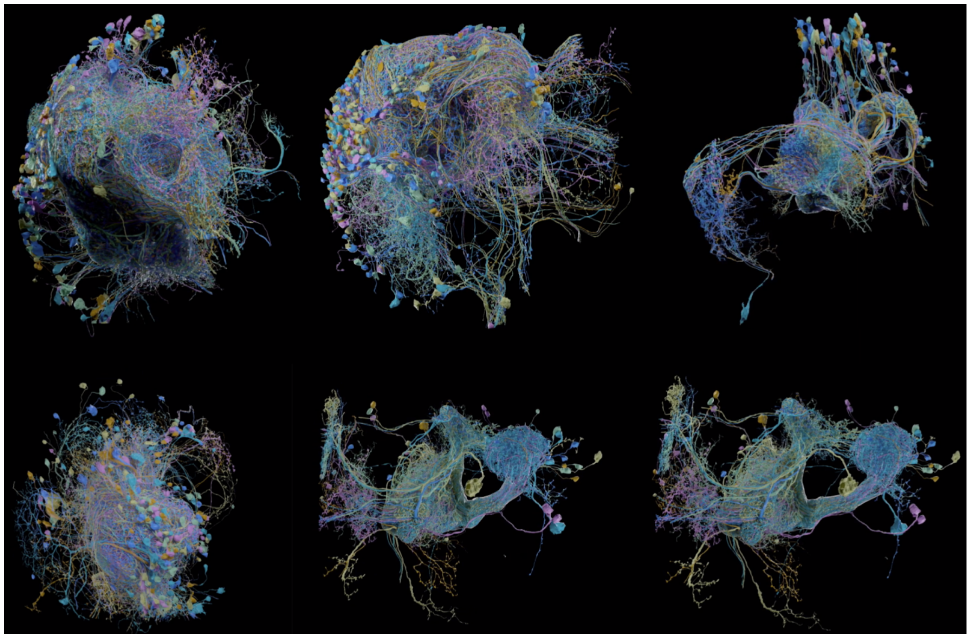

In the summer of 2018, eleven students and I wrote another prospectus entitled, [Amanuensis: The Programmer's Apprentice](), spurred in part by a presentation I gave at the Kavli Futures Symposium on next-generation, open-source neurotechnology in October the previous year. The gist of that talk was the prediction that in the next decade AI technology will enable human-machine symbiotic relationships allowing neuroscientists working with large datasets to interact easily in natural language with AI systems capable of writing code to perform sophisticated data analyses as well as accelerate the development and evaluation of powerful explanatory models. The programmer's apprentice project was resurrected!

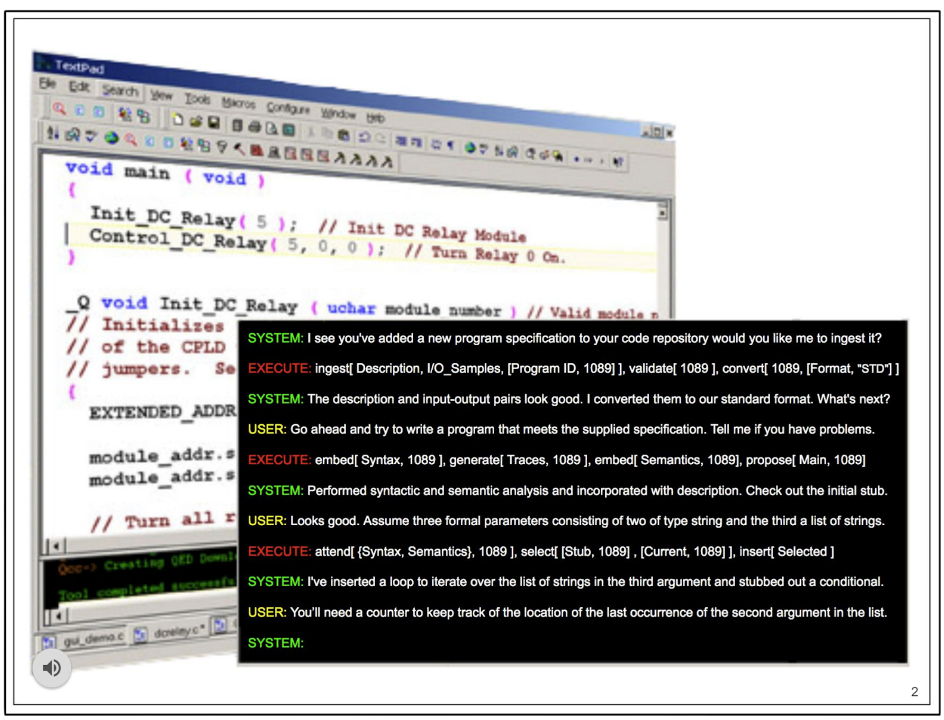

The Programmer's Apprentice was the name of a project started by Charles Rich and Richard Waters at the MIT AI lab in 1987. The goal of the project was to develop a theory of how expert programmers analyze, synthesize, modify, explain, specify, verify, and document programs, and, if possible, implement them. Their research plan was to build prototypes of the apprentice incrementally. Our research plan also involves making incremental steps. However, we will be able to make substantially larger steps by exploiting and contributing to the powerful AI technologies developed during the intervening 30 years with our primary focus on recent advances in applied machine learning and artificial neural networks.

Specifically, we will build on the work of the Interactive Agents Group at DeepMind and borrow from a select set of papers that provide specific technologies relevant to building interactive agents that work closely with humans using natural language to communicate and tackle hard problems that can benefit from the strengths of both human experts and machines with more narrow …

This is an aspirational endeavor. As in the case of our research investigating the prospects for developing advanced technology to spur innovation in neuroscience, the goal here is to examine the prospects for scaling and extending the technology developed by the Interactive Agents Group in order to build systems to assist humans in tackling challenging technical problems.

It is worth pointing out that the general crowd-sourced data collection strategy employed by the IAG team is employed regularly at Google and other technology companies to improve performance on search, machine translation, dialogue management and question answering for Google Assistant. DeepMind's application of this strategy to creating simulated agents that assist in teaching AI ...