Biographical Background: Michael Nielsen

I'm a scientist. I helped pioneer quantum computing and the modern open science movement. I also have a strong side interest in artificial intelligence. All are part of a broader interest in systems and tools that help people discover and create, both individually and collectively.

My interest in quantum computing began in 1992. I am perhaps best known in the field as the author, together with Ike Chuang (MIT), of the standard text on quantum computing. This is the most highly cited work of the past 30 years in physics, and one of the ten most highly cited works ever in physics (based on Google Scholar data available up to ~2015).

I'm especially proud of three research contributions in quantum computing: (1) the fundamental theorem governing the manipulation of entangled quantum states; this kicked off a wide interest in the mathematics of majorization and its relationship to quantum mechanics; (2) a reformulation of quantum computing as a type of geodesic motion in very high dimensional curved spaces; this work is now under intensive study by quantum gravity researchers, who use it to understand black holes; and (3) the discovery and early development of the optical cluster state approach to quantum computing, now being pursued by the company PsiQuantum (latest funding round reported at ~$230 mill).

Other contributions include involvement in the development of quantum gate teleportation, quantum process tomography (used to experimentally characterize quantum gates), and one of the first quantum teleportation experiments, named as one of Science Magazine's Top Ten Breakthroughs of the Year for 1998.

As part of this work, I co-founded and directed the Quantum Information Science Initiative, as Foundation Professor of Quantum Information Science at the University of Queensland. At the time, it was the largest theory-focused quantum computing group in the southern hemisphere, and one of the largest in the world, growing to involve roughly 30 people (faculty, postdocs, students). More broadly, through recruiting, mentorship, and conferences, I helped Australia develop into one of the world's leading countries for work on quantum computing.

While quantum computing is often framed as a promising technology, that wasn't what motivated my interest. I'm fascinated by computers as a means of representing and acting on knowledge, of carrying out processes that we'd call cognition if done by humans. Quantum computers forcefully challenge us to understand the fundamental limits to such processes.

Historically, another line of research has engaged these same questions, albeit from a very different angle. In the 1960s and 1970s, early computing researchers such as Douglas Engelbart, Ivan Sutherland, and Alan Kay began conceiving of computers as tools to augment human cognition. They developed many of the most powerful ideas underlying modern user interfaces, tools extending the human capacity for creativity and discovery.

Inspired by these ideas, in the 1990s I became excited by the promise of the internet to help change the way science is done - through new tools for collaboration, for sharing data and code and ideas, for making meaning in new ways. I saw this promise rapidly realized inside the open source programming community. But it became apparent that many barriers inhibited this within science. Science has developed some strong systems and norms for sharing knowledge (e.g. journal articles) but also many systems that weakly incentivize or outright disincentivize sharing in crucial ways (e.g., data, software and tooling, and the tacit knowledge often crucial in discovery).

Gradually, I developed a belief that the scientific community could develop far better tools, systems, and incentives to support science. In the early 2000s I began to participate in (and helped instigate) a community developing these ideas under the name open science. At its heart, open science is a combination of a pragmatic belief in the value of better tools and systems, and old-fashioned Baconian values: a belief that scientific knowledge is held collectively by humanity, and that science and humanity's interests are best served by a combination of open sharing, collaboration, competition, and robust debate.

I developed this interest as an academic for several years, but in 2008 decided to give up my senior faculty position to work for several years as a full-time advocate for open science. I made this decision because I believed the highest leverage action was to make open science into an issue that every scientist in the world was aware of and had thought seriously about.

You can get some flavour of my work toward this end from my talk at ted.com, in my op-ed in the Wall Street Journal, or in venues like Physics World and Nature. I also wrote a widely-cited book about open science and collective intelligence. It was named one of the best books of 2011 by the Financial Times and by the Boston Globe. (Reviews: Guardian, Financial Times, Nature). Together with work by a small cadre of other people and organizations, this advocacy work helped make open science into a global issue in the scientific community.

By 2015 open science was an established cultural force in science, and seemed likely to continue to gain momentum. I decided to return to full-time research. (I find myself more naturally at home as a researcher than as an advocate; advocacy can slip easily toward uncritical ideology). I obtained unusually independent Research Fellowships, providing support for completely open-ended work. And I turned my attention to work directly on developing novel tools and prototypes to augment human cognition.

This work is difficult to describe briefly, just as work on quantum computing was in the 1990s, and open science in the 2000s: it is early days, and not yet broadly legible. One influential prototype I helped create was a new type of interactive medium known as *Quantum Country*. This looks ostensibly like an educational project, an interactive textbook to teach quantum computing. In fact, it's primarily a research project, introducing many new interface primitives, primitives which change in a fundamental way how users of the medium remember and understand the material. To quote one of the pioneers of modern computing, Alan Kay, the aspirational long-term goal of such projects is "to change the thought patterns of civilization itself".

This is just one of many projects in this vein. Among the other projects is an interactive open textbook about artificial intelligence, which has received more than 2,000 citations, and been accessed by more than 5 million people in 232 countries. And I helped researchers at Google AI conceive and launch an interactive open journal of machine learning, distill.pub, which showcases many new tools, and helps evolve the concept of the scientific journal.