

How Scientific Organizations Adapt to Methodological Advances in Artificial Intelligence: The Impact of AlphaFold1

Gabriel Cavalli

Rotman School of Management, University of Toronto

October 31st, 2022; rev. June 6th, 2023

ABSTRACT

Will Artificial Intelligence (AI) increase or decrease human expertise in science? The purpose of this paper is to examine the impact of a methodological advance in artificial intelligence, specifically the unexpected success of AlphaFold1 in 2018, on the size and expertise composition of scientific organizations. AlphaFold1 represented a significant breakthrough in the scientific subfield of “protein folding” (PF) within computational biology. The unexpected success of AlphaFold1 prompted principal investigators of established academic labs engaged in similar PF-related software development to evaluate whether and how to continue their work in the PF subfield. The findings indicate that the labs affected by AlphaFold1 grew in size, mainly through increased recruitment of PhD students and Postdoctoral researchers. However, principal investigators engaged in varying knowledge-generating strategies for contributing to the PF subfield which reflected their academic backgrounds. The analysis addresses several issues: (i) the approaches through which principal investigators identified new opportunities arising with AlphaFold1, (ii) the impact on competition between academic labs and its consequences for lab composition; and (iii) the implications of AlphaFold1’s computational excellence for the reshaping of projects pursued by human scientists.

Keywords:

Scientific Organizations, Methodological Advances in Artificial Intelligence, Academic Labs, Lab Size, Lab Expertise Composition, AlphaFold1

INTRODUCTION

The size and expertise composition of scientific organizations play a central role in their processes of scientific knowledge production. Past research on this topic has primarily focused on a specific type of scientific organization: teams of co-authoring scientists (Jones, 2009; Uzzi, Mukherjee, Stringer & Jones, 2013; Kaplan & Vakili, 2015; Agrawal, Goldfarb & Teodoridis, 2016; Teodoridis, 2018; Vakili & Kaplan, 2021). Previous studies have revealed that co-authoring teams increase in size when novel ideas become available (Agrawal, Goldfarb & Teodoridis, 2016) or when cost-reducing technologies facilitate collaborations among scientists (Teodoridis, 2018). The utilization of AI-based software may encourage the expansion of scientific teams, as larger teams have the capacity to engage in broader knowledge recombination across different domains (Agrawal, McHale & Oettl, 2018). This effect may be particularly salient in the life sciences, where artificial intelligence is increasingly used for predicting complex biological phenomena (Cockburn, Henderson & Stern, 2019). Furthermore, the breadth of expertise within collaborative teams enables wider knowledge recombinations (Weitzman, 1998; Uzzi et al., 2013; Asgari, Singh & Mitchell, 2017; Vakili & Kaplan, 2021). The depth of expertise among collaborating scientists proves beneficial when technical knowledge is required to identify gaps and anomalies within a scientific domain (Weisberg, 1999; Kaplan & Vakili, 2015).

An important but often overlooked type of scientific organization involved in the production of scientific knowledge is academic labs (Jones, 2021). Common in the life sciences (Sampat & Lichtenberg, 2011; Azoulay, 2019), academic labs are characterized by the presence of a principal investigator who exercises “overarching control” over project execution, recruitment decisions, and task allocation to lab members (Jones, 2021, p. 204). When external advances offer new insights, principal investigators may reconsider their approaches to scientific problems. How do principal investigators of established academic research labs react to unexpected methodological advances?

How do they modify the size and expertise composition of their labs in response to these advances? These questions are at the heart of the research reported here.

This paper examines the impact of a methodological advance in artificial intelligence on the size and expertise composition of established academic labs. Specifically, the paper focuses on the unexpected success of AlphaFold1 in the CASP (Critical Assessment of protein Structure Prediction) initiative in 2018. Held every two years since 1994, the CASP initiative attracts academic labs which compete to develop software for predicting the 3-dimensional (3-D) structure of proteins from their amino acid sequences, a scientific problem known as the “protein folding” (PF) problem. Developed by Google’s DeepMind, AlphaFold1 is an AI-based software program that significantly improved protein structure prediction, marking a breakthrough in the PF subfield of computational biology. As an AI-based software program, AlphaFold1 is not transparent in its inference process. This lack of transparency hinders casual understanding of the process that generates its output, resulting in what is known as the “black box” problem in AI. Additionally, while AlphaFold1 demonstrates remarkable performance in its specific task, it falls short when addressing other protein-related scientific problems.

The response of principal investigators engaged in PF-related software development depends on how they view AlphaFold1. If AlphaFold1 is seen as a nearly definitive solution to the PF problem, then principal investigators may opt to quit the PF research field, since AI-based software has the potential to make human competences and similar software obsolete. If AlphaFold1 is instead viewed as a methodological advance that can inspire further scientific research, principal investigators may choose to adapt to the shock. This paper focuses on the decisions of principal investigators to adapt to AlphaFold1 by examining the changes in the size and expertise composition of their labs. The paper argues that principal investigators engaged in PF-related software development adapted to AlphaFold1 by increasing the size of their labs. Moreover, this paper proposes that principal

investigators decided to either broaden or deepen the expertise composition within their labs to pursue different knowledge-generating strategies. These decisions were based on their academic background and how they anticipated capitalizing on the advances introduced by AlphaFold1.

Overall, the findings suggest that AlphaFold1 opened new areas of inquiry that led to the expansion of established academic labs, and to the recruitment of greater numbers of computational scientists in most PF labs. The exception was among labs run by principal investigators with an interdisciplinary background in both the life sciences and computational sciences. An interdisciplinary background is aligned with a greater emphasis on life sciences after AlphaFold1, suggesting that principal investigators with the most experience attach greater emphasis on the implications of AI for human life science. The results have implications for further research on how artificial intelligence elevates requirements for interdisciplinary insights, scaled computational capabilities, and the organization of science itself.

BACKGROUND AND SETTING

This section first describes the context in which AlphaFold1 was developed (i.e., the CASP initiative and the “protein folding” problem). It then describes what AlphaFold1 represents as a methodological advance in artificial intelligence. Lastly, this section describes how AlphaFold1 is expected to impact established academic labs engaged in PF-related software development.

The CASP initiative and the “protein folding” (PF) problem

The CASP (i.e., Critical Assessment of protein Structure Prediction) initiative is a competition between academic labs which has taken place every two years since 1994.¹ In the CASP initiative, participating academic labs compete by developing software to computationally model a specific scientific phenomenon, i.e., the prediction of proteins’ 3-dimensional (3-D) structures from their

¹ Website: <https://predictioncenter.org/>

amino acid sequences (DeepMind, 2020a). This scientific phenomenon is commonly referred to as the “protein folding” (PF) problem.² The CASP initiative aims “at establishing the current state of the art in protein structure prediction, identifying what progress has been made, and highlighting where future effort may be most productively focused” (Protein Structure Prediction Center, n.d.).³

Over the period from 1972 to 2020, scientists primarily determined the 3-D structures of proteins through experimental physical techniques, such as X-ray crystallography or cryo-electron microscopy. These experimental techniques present numerous problems, including extensive trial error, time consumption, and high costs to determine the 3-D structure of even a single protein (DeepMind, 2020a). Through initiatives such as CASP, scientists develop software as an alternative to experimental physical techniques with the effect of minimizing costs and enabling knowledge production through computational simulations (DeepMind, 2020a).

The metric used in the CASP initiative to compare performance of software developed by competing academic labs is the “Global Distance Test” (GDT), which is a score ranging from 0 to 100, and which is defined as the “percentage of amino acid residues (beads in the protein chain) within a threshold distance from the correct position” (DeepMind, 2020b, “Results from the CASP14 assessment” section, para. 2). The GDT score expresses how well the computational prediction approximates the real 3-D structure of a protein. In the words of Professor John Moult (co-founder and chair of CASP), “On a scale of zero to a hundred, you would expect a GDT over 90 to be a solution to the problem” (DeepMind, 2020c, 2.24-2.33); “a score of around 90 GDT is informally considered to be competitive with results obtained from experimental methods” (DeepMind, 2020d, Introduction section, para. 3).

² For more information on the “protein folding” problem: <https://www.deepmind.com/blog/alphafold-using-ai-for-scientific-discovery-2020>

³ Website: <https://predictioncenter.org/>

DeepMind participated for the first time in the CASP initiative in 2018 with the first version of AlphaFold (called AlphaFold1), which had outperformed the software developed by established PF academic labs in idiosyncratic competitions over the prior 24 years. In December 2018, AlphaFold1 achieved a median score of 58.9 GDT; in December 2020, AlphaFold2 improved drastically, achieving a score of 87.0 GDT (Callaway, 2020; DeepMind, 2020b). This paper considers the unexpected success of AlphaFold1 in CASP in 2018 as a shock (Senior et al., 2020).

What is AlphaFold1?

AlphaFold1 is a publicly available AI-based software program that makes accurate but imperfect inferences about the scientific structure of a specified biological outcome. In the CASP competition, the outcomes are the 3-D structures of pre-specified proteins in their final form (DeepMind, 2020c). AlphaFold1 was developed by DeepMind, a laboratory founded in 2010 and acquired by Google in 2014, which focuses on developing AI-based systems.

AI-based software carries the promise of foundational advancement in science. AlphaFold1 may represent an AI-enabled case of an “invention of a method of inventing” (Griliches, 1957, p. 502; Cockburn, Henderson & Stern, 2019, p. 116). Advances in deep learning are leading to exponential improvements of the efficacy of AI-based software in problems of classification and prediction (Agrawal, Gans & Goldfarb, 2018). AI-based software can be beneficial for problems of classification because it allows faster and more effective “combinatorial search over a vast range of possibilities” (Agrawal, McHale & Oettl, 2022, p. 1), thus allowing recombinations of knowledge across multiple domains (Weitzman, 1998). When it comes to problems of prediction, AI-based software can detect patterns from massive datasets that provide information about biological phenomena (Cockburn, Henderson & Stern, 2019). Thus, AI-based software has the potential to play a transformative role in many scientific endeavors, especially in the life sciences.

However, questions remain about the validity of scientific knowledge produced through AI. AI-based software is quicker and more effective than humans in delivering a solution to problems of classification and prediction (Kurzweil, 2005).⁴ Yet, AI-based software suffers from a “black box” problem, since the process that leads to solutions is updated algorithmically at speeds that are undeciphered and indecipherable to scientists, which makes its action mysterious and construed as unknown (Castelvecchi, 2016). The power of AI-based software in obtaining inferences about scientific phenomena is enabled by the absence of recording of the machine learning that occurs within them. The very success of an AI-based software program such as AlphaFold1 at obtaining a solution to a previously unsolved scientific problem depends on the unobservability of the method of the inference. The computational speed and power of AlphaFold1 is accompanied by this characteristic absence of transparency in how inferences are achieved: “AlphaFold has improved our ability to predict a protein structure from its sequence; but hasn’t directly increased our understanding of how protein sequence relates to structure” (Foldit, 2020, “AlphaFold Performance in CASP” section, para. 4). Like other AI-based software programs, AlphaFold1 integrally does not contribute to causal understanding of how the solution was obtained (Kim, 2022).

As an AI-based software program, AlphaFold1 is relatively narrow in its applicability. It represents an Artificial Narrow Intelligence (ANI), which means that it is very good (although not perfect) at solving a specific and well-defined problem but limited in applicability to other scientific phenomena (Fjelland, 2020). The narrowness of AlphaFold1’s application creates a demand for human expertise in identifying the problems for which the software is appropriate as well as in evaluating its outputs (Athey, Bryan & Gans, 2020), and in assessing the quality of the software itself.

⁴ Ray Kurzweil was hired as Google’s director of engineering in 2012.

In addition to being a publicly available AI-based software program, AlphaFold is multifaceted in nature. It has the potential to become a “General Purpose Technology” (Bresnahan & Trajtenberg, 1995, p. 83; Gambardella & McGahan, 2010, p. 262; Bresnahan, 2023) since it is applicable to basic and commercial uses across different domains, and it is open to improvements and complementarities. Made publicly available by DeepMind on GitHub,⁵ AlphaFold may also represent a public good (Stiglitz, 1999) that has the potential to serve the public interest of the global scientific community. Discerning the nature of AlphaFold is necessary for understanding its potential commercial and scientific impact.

This study focuses on AlphaFold1 as a methodological advance in artificial intelligence, and on its impact on scientific organizations engaged in the development of similar software. Specifically, AlphaFold1 presents an opportunity to delve into the dynamic interplay between organizations and technology (Orlikowski, 1992, 2000) by assessing the ways in which principal investigators of established academic labs engaged in similar software development were compelled to re-assess their own approach to the “protein folding” problem.

The Impact of AlphaFold1 on Principal Investigators of Established Academic Labs

There are a range of possible responses to AlphaFold1 by principal investigators of established academic labs engaged in PF-related software development. If principal investigators view AlphaFold1 as a nearly definitive solution to a scientific problem, then they may see their investments in resources and competences for investigating the “protein folding” problem as made obsolete. Alternatively, they may consider AlphaFold1 as an exceptional methodological advance in artificial intelligence upon which to capitalize, but not as a definitive solution to the “protein folding” problem.

⁵ AlphaFold1: https://github.com/deepmind/deepmind-research/tree/master/alphafold_casp13
AlphaFold2: <https://github.com/deepmind/alphafold>

The methodological advance brought by AlphaFold1 could inspire future research of established academic labs. According to many experts, the “protein folding” problem is “not yet solved” (Moore, Hendrickson, Henderson & Brunger, 2022, p. 507). In the words of Professor Jinbo Xu (Associate Professor at the Toyota Technological Institute at Chicago, and Senior Fellow of the Computation Institute at the University of Chicago in Illinois), “The open-source nature of the tools means that the scientific community should be able to build on the advances to create even more powerful and useful software” (Callaway, 2021, Introduction section, para. 4). According to experts such as Professor Janet M. Thornton (Director Emeritus at the European Bioinformatics Institute) and Professor David T. Jones (Professor in Bioinformatics at the University College of London), AlphaFold will stimulate “fierce competition to improve the method even further and new applications of machine learning to help illuminate proteomes and their many interactions” (Jones & Thornton, 2022, Abstract section, p. 15). After the release of the source code, PF labs have started re-interpreting and re-producing AlphaFold1’s (and AlphaFold2’s) code with their own twists.⁶

Therefore, AlphaFold1 may not represent a nearly definitive solution to a scientific problem that leads established researchers to stop investing. Rather, DeepMind’s participation in PF may represent the beginning of a period of increased dynamism in a scientific subfield. To pursue the opportunities revealed by AlphaFold1, principal investigators of established academic labs may have responded by recombining their organizations (Levinthal, 1997).

THEORY

The Theory section begins by considering the decision faced by principal investigators of established academic labs: either to quit the impacted subfield or to adapt to the methodological advance brought by AlphaFold1. Then, this section proposes hypotheses about the impact of AlphaFold1 on the size

⁶ After AlphaFold1: <https://www.biorxiv.org/content/10.1101/830273v2.abstract>
After AlphaFold2: https://github.com/biolists/folding_tools

of established academic research engaged in PF-related software development. Lastly, the theoretical discussion formulates hypotheses about the impact of AlphaFold1 on principal investigators' choices in terms of the expertise composition of their labs, contingent upon their academic background.

Decisions by Principal Investigators of Established Academic Labs: Quit or Adapt?

Academic research in the computational life sciences is predominantly conducted in labs managed by principal investigators (Sampat & Lichtenberg, 2011; Azoulay, 2019). Decisions about which scientific projects to undertake or abandon are made primarily by principal investigators (Jones, 2021). In response to the shock represented by AlphaFold1, principal investigators may thus react in one of two ways: they may decide to quit the PF scientific subfield; or they may instead decide to stay and adapt to the methodological advance brought by AlphaFold1 (Figure 1).

Insert Figure 1 about here

Two main reasons explain why principal investigators may decide to quit the PF subfield. First, by automating tasks, AI-based software may make human competences obsolete (Acemoglu & Restrepo, 2019) and similar software useless. This reason was summarized by Professor Mohammed AlQuraishi (Principal Investigator of the AlQuraishi Laboratory and Assistant Professor in the Department of Systems Biology at Columbia University) in 2018:

“In a delicious twist of irony, we the people who have bet their careers on trying to obsolete crystallographers are now worried about getting obsoleted ourselves. [...] [F]or some of us it will make sense to go into industrial labs, while for others it will mean staying in academia but shifting to entirely new problems or structure-proximal problems that avoid head-on competition with DeepMind” (AlQuraishi, 2018, “What just happened?” section, para. 1 and 3).

Second, when an innovation solves a problem definitively, no further research on that problem is needed. After a scientific problem is solved, labs shift their efforts to areas in which there are greater opportunities, with “unmet therapeutic needs and unexploited biological mechanisms” (Pammolli, Magazzini & Riccaboni, 2011, p. 428). Moreover, funding and talent tend to flow towards academic labs and research fields that have higher likelihoods of new impactful discoveries (Merton, 1968; Latour & Woolgar, 1986; Kitcher, 1990; Strevens, 2003).

Current literature documents the potential complementarity between humans and machines in decision-making processes (Puranam, 2021; Agrawal, Gans & Goldfarb, 2022). If AlphaFold is viewed as a methodological advance in artificial intelligence that can benefit established academic labs, principal investigators may elect to adapt to the shock. Scientists make significant investments of time, effort, and resources to enter a knowledge domain (Fleming, 2001; Jones, 2009; Leahey, Beckman & Stanko, 2017). If a scientific subfield is impacted by a novel methodological advance and scientists expect to build upon it, principal investigators may choose to adapt rather than quit. This assertion aligns with prior theoretical research demonstrating that organizations are unlikely to give up on past efforts (Perignat & Fleming, 2022) due to sunk costs (Stiglitz, McFadden & Peltzman, 1987), and the potential value of accumulated expertise.

This paper focuses on the decision by principal investigators of established academic labs engaged in PF-related software development to adapt. Conditional on staying in the impacted research subfield, the paper examines how they responded to the shock in terms of size and expertise composition of their labs.

Decision by Principal Investigators on Lab Size

As ideas are “getting harder to find” (Bloom, Jones, Van Reenen & Webb, 2020, p. 1104) and a growing “burden of knowledge” is placed upon scientists’ shoulders (Jones, 2009, p. 283), teams

performing scientific research are becoming bigger and more interdisciplinary (Haeussler & Sauermann, 2020). By growing in size, scientific organizations can improve their ability to adapt and capitalize on scientific breakthroughs, such as methodological advances in artificial intelligence.

Past literature has focused on some of the factors that lead to increases in size of scientific organizations. One is the need to gather new competences in response to ground-breaking technological advances (Lacetera, Cockburn & Henderson, 2004). Adoption of science-based drug discovery by pharmaceutical firms was shown to be correlated with the hiring of star scientists – a finding consistent with the idea that changes in organizations’ capabilities can be obtained through recruitment processes (Lacetera, Cockburn & Henderson, 2004). A second factor shown to lead to increases in size of scientific organizations is the influx of novel ideas after institutional shocks (Agrawal, Goldfarb & Teodoridis, 2016). The fall of the Soviet Union led to an inflow of Soviet mathematicians into Western universities (Agrawal, Goldfarb & Teodoridis, 2016). In subfields of mathematics in which Soviet mathematicians were strongest, teams of co-authoring scientists grew disproportionately in size (Agrawal, Goldfarb & Teodoridis, 2016). Third, novel platforms which reduce costs for integrating diverse competences have been shown to stimulate collaborations among scientists (Teodoridis, 2018). The hacking of Microsoft Kinect reduced entry costs for generalists, who could access the “motion sensing” research area and establish collaborations with outside-area specialists (Teodoridis, 2018). Novel technological advances can thus open new opportunities for knowledge recombinations (Teodoridis, 2018). Moreover, research in the field of strategic management has addressed the role played by AI-based software on the size of scientific teams. Current literature argues that AI-based software encourages increases in scientific teams’ size because larger teams can recombine more widely across knowledge domains (Agrawal, McHale & Oettl, 2018, 2022). These findings suggest that increases in lab size are beneficial because knowledge

resulting from multiple scientists can yield more impactful discoveries (Wuchty et al., 2007). Thus, the paper proposes the following hypothesis:

***Hypothesis 1.** After a methodological advance in artificial intelligence appears within a scientific subfield, exposed principal investigators are likely to adapt by increasing the size of their labs.*

Principal Investigators' Academic Background as a Contingency for Changes in Lab Expertise Composition

Decisions made by principal investigators about the expertise composition of their labs are influenced by various factors. Resource allocation significantly impacts the types of expertise present in a lab (Zhuo, 2022) and the knowledge-generating strategies pursued by scientists (Boudreau, Guinan, Lakhani & Riedl, 2016). While some principal investigators may allocate resources to recruit expertise for the exploitation of existing scientific opportunities, others may instead opt for more exploratory and risky projects in lesser-known research subfields (Zhuo, 2022). Another factor that can affect the expertise composition and knowledge-generating strategies of academic labs is the cognitive perspective of key organizational decision-makers (Kaplan & Tripsas, 2008), manifested through principal investigators' academic background in this setting. This paper explores the role of principal investigators' academic background on decisions regarding the expertise composition within their labs. The aim of the paper is to gain insights into how different academic backgrounds have led principal investigators to pursue distinct knowledge-generating strategies through decisions on lab expertise composition, in response to the appearance of a methodological advance in artificial intelligence within their scientific subfield.

According to past literature, innovation can be the result of both broad recombinative processes (Weitzman, 1998; Fleming, 2001) and deep expertise within a knowledge domain (Weisberg, 1999; Taylor & Greve, 2006; Kaplan & Vakili, 2015). Greater breadth of knowledge

inputs through expanded expertise areas allows for more novel knowledge recombinations, which in turn lead to more impactful innovations (Weitzman, 1998; Schilling & Green, 2011; Uzzi et al., 2013). Expertise depth has been shown to stimulate innovation in domains requiring deep knowledge for identifying foundational gaps and anomalies (Weisberg, 1999; Taylor & Greve, 2006; Kaplan & Vakili, 2015). This paper proposes that principal investigators decided to either broaden or deepen the expertise composition within their labs to pursue different knowledge-generating strategies. These decisions were based on their academic background and how they anticipated capitalizing on the advances introduced by AlphaFold1.

If the knowledge-generating strategy pursued by principal investigators engaged in PF focuses on the optimization of software efficacy, and if they expect to be able to compete with AlphaFold1 by incorporating its advances into their own proprietary software, then it makes sense for them to deepen computational expertise within their own labs. Scientists with deep expertise are more likely to absorb and use knowledge at the frontier (Teodoridis, Bikard & Vakili, 2019). Principal investigators who possess an academic background in the computational sciences may thus decide to deepen the expertise of their labs through recruitment of individuals with a similar background, especially if they have previous experience in adopting and absorbing valuable code (King & Lakhani, 2011). Knowledge similarity between principal investigators and newly recruited lab members would ensure more effective knowledge production in the form of more effective PF-related software (Lane, Ganguli, Gaule, Guinan & Lakhani, 2021). Therefore, when a methodological advance in artificial intelligence appears within a scientific subfield, principal investigators who hold the belief that they can compete with AlphaFold1 may respond by deepening expertise within their labs. Principal investigators with an academic background in the computational sciences may seek to absorb frontier knowledge as effectively as possible through recruitment of lab members with a similar background.

Hypothesis 2a. *After a methodological advance in artificial intelligence appears within a scientific subfield, exposed principal investigators are likely to deepen the expertise within their labs if their academic background leads them to believe that they can incorporate the novel knowledge brought by the advance.*

If the knowledge-generating strategy pursued by principal investigators engaged in PF focuses on the understanding of biological mechanisms, they may find useful to incorporate within their labs novel knowledge components to strengthen their computer-modelling competences. In the attempt to improve the computational competences of their labs, principal investigators with an academic background in life sciences may be willing to broaden the expertise of their labs. Expertise breadth can help scientists and organizations maximize their innovative output (March, 1991; Jeppesen & Lakhani, 2010; Leiponen & Helfat, 2011), especially in broadly applied and modular technological domains such as software development (Vakili & Kaplan, 2021). Therefore, when a methodological advance in artificial intelligence appears within a scientific subfield, principal investigators who focus on the wider implications of “protein folding” may respond by broadening expertise with their labs. Principal investigators with an academic background in the life sciences may seek to integrate as many diverse knowledge components as possible and strengthen their labs’ software development competences through recruitment of lab members with a different background.

Hypothesis 2b. *After a methodological advance in artificial intelligence appears within a scientific subfield, exposed principal investigators are likely to broaden the expertise within their labs if their academic background leads them to believe that they need additional knowledge components to benefit from the advance.*

If the knowledge-generating strategy pursued by principal investigators engaged in PF focuses on the tractability of biological phenomena, they may consider a methodological advance such as AlphaFold1 as a source of ideas for research on other PF-related scientific problems. When competition for creativity reaches excessive intensity, innovators may stop investing (Boudreau, Lakhani & Menietti, 2016; Gross, 2020), especially in cases where the problem at hand has become less uncertain (Boudreau, Lacetera & Lakhani, 2011). Principal investigators with a background in both the life and the computational sciences may believe that the competition between academic labs has become too intense and decide to investigate novel PF-related problems which cannot be addressed through AlphaFold1. To identify novel PF-related problems, principal investigators with an interdisciplinary background may decide to reduce over-specialization and over-reliance on past computational approaches (Bryan, 2017), and attach greater emphasis on expertise related to the experimental life sciences. Therefore, when a methodological advance in artificial intelligence appears within a scientific subfield, principal investigators who believe that methodological advances in artificial intelligence are limited in their ability to tract different biological phenomena may respond by reshaping the expertise composition of their labs through lab rotations. Principal investigators with an interdisciplinary academic background may seek to tackle novel PF-related problems by retaining lab members who have the expertise to identify these novel problems.

***Hypothesis 2c.** After a methodological advance in artificial intelligence appears within a scientific subfield, exposed principal investigators are likely to reshape the expertise composition within their labs if their academic background leads them to believe that they can extend the novel knowledge brought by the advance to related scientific problems.*

DATA AND EMPIRICAL ANALYSIS

This section describes how data about labs, lab members, and principal investigators were collected. Then, it describes how the variables of interest (size, by role and expertise type) were constructed. Lastly, this section describes the estimation strategy for measuring how established academic labs engaged in PF-related software development changed their size and expertise composition after the unexpected success of AlphaFold1.

Data Collection

The analysis considers data from 177 academic labs, which are based in 120 institutions located in 25 countries (as shown in [Figure 2](#) and [Figure 3](#)). [Figure 4](#) illustrates the distribution of institutions by country and scientific subfield. These 177 labs were identified by searching through the websites of academic institutions worldwide, specifically in the departments of Biology, Chemistry, Biochemistry, Medicine, Physics, Computer Science, Mathematics, and Statistics. Each lab has its own website.

Insert Figures 2, 3 and 4 about here

All the labs considered in the analysis are engaged in software development for computationally modelling protein-related phenomena (i.e., protein folding, protein-protein interactions, protein-DNA docking, etc.). In the scientific field of protein-related computational biology, the main output of academic labs is software. To test the validity of their software, principal investigators participate in various community-wide initiatives, in which they compete against each other. Examples of these initiatives are: (i) the Critical Assessment of protein Structure Prediction

(CASP);⁷ (ii) the Critical Assessment of protein Function Annotation algorithms (CAFA);⁸ (iii) the Critical Assessment of Genome Interpretation (CAGI);⁹ (iv) the Critical Assessment of Metagenome Interpretation (CAMI);¹⁰ (v) the Critical Assessment of PRedicted Interactions (CAPRI);¹¹ (vi) the Continuous Automated Model EvaluatiOn (CAMEO);¹² (vii) the Critical Assessment of Intrinsic protein Disorder (CAID);¹³ and others. When participating in these initiatives, principal investigators test the efficacy of their own software against the efficacy of software developed in other academic labs. After participating in these initiatives, principal investigators make their software public on their labs' websites.

To determine if an academic lab is engaged in PF-related software development, the analysis considers whether it has participated in the Critical Assessment of protein Structure Prediction (CASP) initiative at least once between 2014 and 2018, prior to the unexpected success of AlphaFold1. The participation of principal investigators in the CASP initiative during this period was determined by extracting their names from the “Abstract Books” of each edition. These “Abstract Books” provide information on the labs that participated in each edition of the CASP initiative.¹⁴ Participation in the CASP initiative is used as a proxy to operationalize whether an academic lab is engaged in PF or not, and thus whether a lab is classified as treated or not. The limitations of this approach will be described in the Discussion section.

⁷ Website of CASP: <https://predictioncenter.org/>

⁸ Website of CAFA: <https://www.biofunctionprediction.org/cafa/>

⁹ Website of CAGI: <https://genomeinterpretation.org/>

¹⁰ Website of CAMI: <https://data.cami-challenge.org/>

¹¹ Website of CAPRI: <https://www.capri-docking.org/>

¹² Website of CAMEO: <https://www.cameo3d.org/>

¹³ Website of CAID: <https://idpcentral.org/caid>

¹⁴ Labs participating in the CASP initiative in 2018: https://predictioncenter.org/casp13/doc/CASP13_Abstracts.pdf

Labs participating in the CASP initiative in 2016: https://predictioncenter.org/casp12/doc/CASP12_Abstracts.pdf

Labs participating in the CASP initiative in 2024: https://predictioncenter.org/casp11/doc/CASP11_Abstracts.pdf

When considering all labs, the analysis includes 55 treated labs and 122 control labs, resulting in a total of 177 labs. When focusing only on senior labs (i.e., labs that have been operational throughout the entire 2015-2021 period), the analysis includes 42 treated labs and 101 control labs, resulting in a total of 143 senior labs (Figure 5).

Insert Figure 5 about here

Data about the size and expertise composition of the labs were collected manually using the following approach. The names and surnames of current and former lab members were gathered from the labs' websites. Many labs provide information about both current members and former members (often referred to as labs' alumni), including their roles in the lab (such as Research Assistant, PhD Student, Postdoctoral researcher, Staff, Visiting Professor, etc.) and the duration of their time in the lab (start and end dates)¹⁵. In cases where websites did not provide information about former members or their period in the lab, historical versions of the websites were accessed using Wayback Machine¹⁶, an Internet archive that stores snapshots of websites over time. Wayback Machine was also used to verify the information obtained from current websites.

After collecting the names and surnames of current and former lab members, information about each individual was obtained from LinkedIn. On LinkedIn, data regarding the period spent in the lab was gathered and validated against the information already provided on the labs' websites. Additionally, the complete academic background of each lab member, including the fields in which they obtained their Bachelor's, Master's, and/or PhD degrees, was recorded.

¹⁵ An example: <https://digbio.missouri.edu/our-team/>

¹⁶ Wayback Machine: <https://web.archive.org/>

An example illustrates how the dataset was built. On the website of Lab k , an individual was reported to have worked as a Research Assistant from 2014 to 2017. After collecting this information, her name and surname were searched on LinkedIn. This individual reported to have obtained a Bachelor's degree in Biology in the period from 2012 to 2015; a Master's degree in Computer Science in the period from 2015 to 2018; and a PhD in Computational Biochemistry in the period from 2018 to 2022. Moreover, in her professional experience on LinkedIn, this individual confirmed to have worked as a Research Assistant in Lab k from 2014 to 2017. In the dataset, this individual was coded as follows: in 2014, as an Undergraduate Research Assistant with expertise in Biology; in years 2015, 2016, and 2017, as a Master's Research Assistant with expertise in Biology and Computer Science. Her experience in Computational Biochemistry was not considered because it arose after she left the lab.

The analysis considers information about labs' size and expertise composition during a seven-year period from 2015 to 2021 (inclusive), which includes four years prior to AlphaFold1 in 2018, and three years subsequent to it.

Variables of Interest: Lab Size and Expertise Composition

Lab size is measured as the yearly count of unique lab members. Expertise composition within labs is measured using information about the academic backgrounds of each lab member. Lab members, including principal investigators, were classified into three main expertise types based on their academic backgrounds: (i) Life Sciences (LS); Computational Sciences (CS); and (iii) a combination of both Life and Computational Sciences (LS+CS).

Lab members were classified as having a Life Sciences (LS) academic background if they obtained their degrees exclusively in disciplines related to experimental life sciences, such as Biology, Chemistry, Biochemistry, Medicine, Pharmacology, Genetics, or Microbiology. Lab members with a

Computational Sciences (CS) academic background obtained their degrees in fields related to computational sciences, such as Computer Science, as well as other sciences not directly associated with experimental life sciences such as Mathematics, Physics, Statistics, or Engineering. Lab members with an academic background in both Life and Computational Sciences (LS+CS) possess expertise in both domains. For example, they may have obtained their degrees in disciplines which are the intersection of both the life and computational sciences, such as Bioinformatics, Computational Biochemistry, or Biophysics. LS+CS members may also have obtained distinct degrees in the two separate types of sciences throughout their academic careers. For instance, they may have obtained a Bachelor's degree in Biology and a PhD degree in Computer Science.

Expertise composition is calculated as the number of unique lab members in the LS, CS, and LS+CS groups. Changes in the expertise composition of labs are measured by considering the alignment between the academic backgrounds of principal investigators and recruited lab members. For instance, if a principal investigator with a CS academic background increases recruitment of profiles with a CS background, this indicates a deepening in expertise within the lab; if the same CS principal investigators increases recruitment of profiles a LS and/or LS+CS backgrounds, this indicates a broadening in expertise within the lab.

Finally, lab members are classified into four primary roles within the lab: (i) Undergraduates and Master's students; (ii) PhD students and Postdoctoral researchers; (iii) Staff; and (iv) Professors.

Estimation Strategy

Following a methodological advance in artificial intelligence in the PF subfield, this paper investigates the subsequent changes in the size and expertise composition of established academic labs engaged in PF-related software development. To examine these changes, the analysis utilizes a panel dataset with Lab k as the unit of analysis. Specifically, the study assesses differences between labs engaged in PF-related software development and labs engaged in software development for

computationally modelling other protein-related phenomena, before and after AlphaFold1. First, the analysis estimates the following two-way fixed effects equation (Teodoridis, 2018):

$$LabSize_{k,t} = \beta_0 + \beta_1 \cdot PostAlphaFold1_t \cdot LabProteinFolding_k + \theta_k + \gamma_t + \epsilon_{kt}$$

Where $LabSize_{k,t}$ represents the number of unique members in Lab k in year t . $PostAlphaFold1_t$ is an indicator variable that takes the value of 1 for years after 2018 (i.e., 2019, 2020 and 2021). $LabProteinFolding_k$ is an indicator variable that takes the value of 1 if Lab k has participated in CASP at least once between 2014 and 2018, and 0 otherwise. θ_k and γ_t are lab- and year-fixed effects, respectively.

To estimate the impact of principal investigators' academic background on changes in lab size and expertise composition, the analysis employs the following triple difference specification (Moser & Voena, 2012):

$$\begin{aligned} LabSize_{i,k,t} = & \beta_0 + \beta_1 \cdot ExpertiseType_i + \beta_{2t} \cdot LabProteinFolding_k \cdot YearPostAlphaFoldV1_t \\ & + \beta_3 \cdot ExpertiseType_i \cdot LabProteinFolding_k + \beta_{4t} \cdot ExpertiseType_i \cdot YearPostAlphaFoldV1_t \\ & + \beta_t \cdot ExpertiseType_i \cdot LabProteinFolding_k \cdot YearPostAlphaFoldV1_t + \theta_k + \gamma_t + \epsilon_{kt} \end{aligned}$$

Where $LabSize_{i,k,t}$ represents the number of unique members in Lab k led by a principal investigator with $ExpertiseType$ i in year t . $ExpertiseType$ i corresponds to the three main academic backgrounds (LS, CS, and LS+CS).

The equations are estimated using a dataset covering the period from 2015 to 2021, which includes four years before the unexpected success of AlphaFold1 in 2018, and three years after. The dataset is structured at the lab-year level.

RESULTS

Changes in Labs' Size

Table 1 provides results about the changes in the size of labs engaged in PF-related software development and compares these changes with labs that develop software for computationally modelling other protein-related phenomena. To assess these changes, Poisson Pseudo-Likelihood regressions were conducted, since the dependent variable involves yearly counts of lab members (Silva & Tenreyro, 2006).

Table 1 includes multiple columns, focusing on specific roles, including all labs and senior labs (i.e., labs that have been operational throughout the entire 2015-2021 period). Columns 1 and 2 examine all lab members collectively. Columns 3 and 4 examine the recruitment of Undergraduate and Master's students. The results for PhD students and Postdoctoral researchers are presented in Columns 5 and 6. Columns 7 and 8 present results about the recruitment of Staff members, while Columns 9 and 10 present results related to the Professors.

Insert Table 1 about here

The coefficient of the interaction term ($PostAlphaFoldV1_t * Lab_ProteinFolding_k$) is positive and statistically significant when PhD students and Postdoctoral researchers are considered (Columns 5 and 6 of Table 1). It is not statistically significant for the other categories of lab members. Since all labs hire PhD students and Postdoctoral researchers, but not necessarily the other categories, it is reasonable to focus exclusively on examining the effects on PhD students and Postdoctoral researchers.

The results validate Hypothesis 1, demonstrating that following the unexpected success of AlphaFold1, principal investigators engaged in PF-related software development have significantly

increased their recruitment of PhD students and Postdocs compared to principal investigators of labs focused on software development for other protein-related phenomena. This finding supports the hypothesis that the methodological advance represented by AlphaFold1 has motivated principal investigators of PF labs to expand their lab sizes by recruiting more PhD students and Postdoctoral researchers.

To check for pre-trends, the analysis presents additional figures. Figure 6 provides descriptive evidence showing a disproportionate increase in the average size of PF labs after 2018. When all PF labs are considered (blue line in Figure 6), the average number of PhD students and Postdoctoral researchers rose from approximately 7 during the period from 2015 to 2018, to around 9 in 2021.

Insert Figure 6 about here

Figures 7a, 7b, 7c, 7d present the estimated yearly differences across different roles within all labs; Figures 8a, 8b, 8c and 8d focus on senior labs. Examining Figures 7b and 8b, there are no significant differences in lab size concerning the recruitment of PhD students and Postdoctoral researchers until the success of AlphaFold1 in 2018. After 2018, the difference in size between PF labs and labs engaged in software development for other protein-related phenomena starts to widen.

Insert Figures 7a, 7b, 7c, 7d about here

Insert Figures 8a, 8b, 8c, 8d about here

To examine the influence of principal investigators' academic background on their response to AlphaFold1, a triple difference estimation specification is employed. Figure 9a shows that principal

investigators with a CS background increased their recruitment of PhD students and Postdocs. [Figure 9b](#) indicates that principal investigators with an LS background did not significantly alter the size of their labs. Interestingly, [Figure 9c](#) reveals that principal investigators with a LS+CS background reduced the number of PhD students and Postdoctoral Researchers. These findings remain consistent when considering only senior labs (i.e., labs that have been operational throughout the entire 2015-2021 period), as depicted in [Figure 10a, 10b, 10c](#). Principal investigators with different academic backgrounds showed distinct responses to the methodological advance represented by AlphaFold1. The next section will delve into how these distinct responses translated into different knowledge-generating strategies.

Insert Figures 9a, 9b, 9c about here

Insert Figures 10a, 10b, 10c about here

Changes in Lab Expertise Composition, Contingent upon Principal Investigators’ Academic Background

To test Hypotheses 2a, 2b and 2c and explore how principal investigators with different academic backgrounds engaged in distinct knowledge-generating strategies through decisions on lab expertise composition, a triple difference estimation specification is employed here as well. Also lab members (i.e., PhD students and Postdoctoral researchers) are categorized into the three main expertise types: LS, CS, and LS+CS.

[Figure 11a](#) illustrates that principal investigators with a CS background increased their recruitment of PhD students and Postdoctoral researchers with a CS background. [Figure 11b](#) shows

that CS principal investigators stopped recruiting profiles with an LS background. Figure 11c indicates no significant changes in the recruitment of LS+CS lab members. These findings validate Hypothesis 3a, suggesting that principal investigators with a CS academic background, who have the greatest potential to benefit from a methodological advance in artificial intelligence, primarily adapted to AlphaFold1 by deepening the expertise within their labs and recruiting members with a similar CS background. One possible explanation is that these principal investigators are more likely to incorporate and utilize cutting-edge knowledge, believing they can leverage AlphaFold1's advancements to enhance their own proprietary software.

Insert Figures 11a, 11b, 11c about here

Figure 12a illustrates that principal investigators with an LS background increased their recruitment of PhD students and Postdoctoral researchers with a CS background. Figure 12b shows that LS principal investigators stopped recruiting profiles with an LS background. Figure 12c indicates no significant changes in the recruitment of LS+CS lab members. These findings validate Hypothesis 3b, suggesting that principal investigators with an LS academic background adapted to AlphaFold1 by broadening the expertise within their labs and recruiting profiles with a different (CS) background. One possible explanation is that LS principal investigators are expanding the capabilities of their labs by strengthening software competences in order to better understand the broader implications of protein folding.

Insert Figures 12a, 12b, 12c about here

Lastly, Figure 13a shows that principal investigators with an LS+CS background reduced their recruitment of PhD students and Postdoctoral researchers with a CS background. Moreover, Figure 13b shows that LS+CS principal investigators stopped dismissing profiles with an LS background. Also here, Figure 12c indicates no significant changes in the recruitment of LS+CS lab members. These findings validate Hypothesis 3c, suggesting that principal investigators with an LS+CS background adapted to AlphaFold1 by reshaping the expertise composition of their labs, dismissing profiles who are over-specialized in building code (i.e., CS PhD students and Postdoctoral researchers), while also retaining profiles specialized in identifying new PF-related scientific problems to target (i.e., LS profiles).

Insert Figures 13a, 13b, 13c about here

DISCUSSION

This paper examines the impact of AlphaFold1, a methodological advance in artificial intelligence, on the size and expertise composition of established academic labs engaged in PF-related software development. Results show that in response to AlphaFold1 PF labs became larger, and that the way in which principal investigators modified the expertise composition of their labs depended on their academic background.

The results are not without limitations, although each is somewhat mitigated by the features of the setting. One fundamental limitation is represented by the nature of the shock considered. Google's DeepMind launched two versions of AlphaFold: AlphaFold1 in 2018, and AlphaFold2 in 2020. Causal claims regarding the impact of these methodological advances on the size and expertise composition of established academic labs may confound the effect of the two versions of AlphaFold. This paper treats AlphaFold1 as the main shock, because it was immediately recognized as an

unexpected event of an industrial lab entering the PF subfield, significantly impacting established PF academic labs (AlQuraishi, 2019). The accuracy score in 2018 (58.9 GDT) associated with AlphaFold1 suggests that its impact was sufficiently strong to be felt by principal investigators engaged in PF-related software development; but not strong enough to have an impact on labs engaged in software development for computationally modelling other protein-related phenomena (Abriata, 2021). However, the nature of the shock does not allow to disentangle the impact of AlphaFold1 from the impact of AlphaFold2, especially after 2020.

A second limitation refers to the criterion used for identifying “treated” labs (i.e., academic labs that have participated in CASP at least once in the period 2014-2018). Ideally, treatment indicators should be built through measures of technological proximity (Bloom, Schankerman & Van Reenen, 2013). In other words, also labs that have not participated in the CASP initiative are likely to have developed some PF-related software throughout their history. Moreover, some labs in the control group may have participated in the CASP initiative before 2014. Building measures of labs’ technological proximity to the PF subfield would allow a more nuanced understanding of labs’ response to AlphaFold1. The present analysis uses participation in the CASP initiative at least once between 2014 and 2018 for two reasons. First, the paper hypothesizes that principal investigators who have participated in the CASP initiative in the same period in which DeepMind joined are those most likely to respond to the shock. Second, the CASP initiative is considered the “Olympics for protein folding” (Toews, 2021, “A Triumph of AI” section, para. 1), in which participate academic labs from all over the world. The analysis assumes that all labs able to participate in the CASP initiative do participate.

One may argue that sample selection issues might bias the results. It is possible that only labs located in the most resourceful institutions can participate in the “Olympics for protein folding” (Toews, 2021). Principal investigators participating in the CASP initiative would thus be the ones

most likely to act on the size and expertise composition of their labs in response to a methodological advance, while principal investigators not participating in CASP would not have sufficient resources to do so. This issue may lead to an over-estimation of AlphaFold's impact on lab size and expertise composition. To address this concern, Figure 4 shows distribution of institutions by country and scientific subfield. Although sample selection issues may still affect the results, Figure 4 shows that PF labs are not located in the most resourceful institutions, nor are they disproportionately located in the US.

Another limitation refers to how the dataset and variables were built. Measures of expertise composition were built by considering the entire academic background of each lab member. In the approach used in this paper, equal weight was assigned to the competences gained by lab members during their Bachelor's, Master's and/or PhD programs. However, competences may decay over time. Lab members' last academic achievements are likely to be the most reliable indicators of their expertise. Assigning greater weight to competences gained later (and lower weight to competences acquired in the earlier stages of a lab member's academic career) would allow a more nuanced understanding of principal investigators' response. Further research on this point is warranted.

The present study considers a very specific setting and an unusual shock. The analysis considers a subfield of computational biology, in which principal investigators of labs engaged in PF-related software development adapt in response to a methodological advance in artificial intelligence which may make their work obsolete. Questions about the generalizability of this setting and the results obtained necessarily arise. Two features of this study may extend the validity of the theoretical framework proposed, and of the findings obtained. First, the analysis considers academic labs that are: (i) headquartered in 25 countries all over the world; (ii) active for the solution of similar scientific problems; (iii) apply the same (scientific) method in institutions with similar purposes (i.e., academic institutions); but (iv) have discretion over which competences to recruit. This feature allows to isolate

the decisions of key organizational decision-makers (in this context, principal investigators), while holding constant cultural, technological, and institutional factors. Second, the analysis investigates how novel and more powerful AI-based software impacts organizations in a scientific subfield. In a world characterized by increasing competition between Big Tech companies for developing the most powerful algorithms, studies about the impact of a new AI-based software program on decisions of established organizations about human capital and organizational composition can provide useful insights to practitioners and policymakers.

CONCLUSION

This paper investigates the impact of a methodological advance in artificial intelligence on the size and expertise composition of established academic labs. The unexpected success of AlphaFold1 (an AI-based software program for predicting proteins' 3-D structure) in 2018 was an event that significantly impacted the work of established academic labs engaged in similar PF-related software development. Results show that impacted principal investigators hired more PhD students and Postdoctoral researchers. However, their response varied depending on their academic background: (i) principal investigators with a CS academic background adapted to AlphaFold1 by deepening the expertise within their labs, recruiting more CS PhD students and Postdocs; (ii) principal investigators with a LS academic background adapted by broadening the expertise of their labs, recruiting more CS PhD students and Postdocs, and stopping recruitment of LS profiles; (iii) principal investigators with an LS+CS adapted by reducing the number of CS lab members while retaining LS profiles. As noticed by observers at the 15th edition of the CASP initiative in 2022 (four years after AlphaFold1, and two years after AlphaFold2), “the most successful teams were those that had adapted and built on AlphaFold in various ways” (Callaway, 2022, “Matchmaking” section, para. 5).

Concerns have arisen regarding the negative consequences of artificial intelligence, and the rising levels of unemployment induced by job automation (Korinek & Stiglitz, 2018). Theoretical models have shown that artificial intelligence is likely to automate low-skilled tasks (Acemoglu & Restrepo, 2019). This paper proposes an alternative explanation of the interaction between humans and artificial intelligence. Artificial intelligence has the potential to strongly augment human capacities (Athey, Bryan & Gans, 2020). Past literature has focused on artificial intelligence's ability to produce outcomes; however, artificial intelligence, by its own design, cannot provide nor grasp the abstract cause of phenomena. Artificial intelligence does act in a way to modify internal parameters, in order to maximize an objective function (LeCun, Bengio & Hinton, 2015). But it does not generate understanding of the abstract cause of phenomena – an ability coded into machines by their builders (Turing, 1950). Nor can artificial intelligence answer why it performs a task – this feature explains the intrinsic impossibility for machines to be held accountable for their actions (Galasso & Luo, 2018). This paper is aligned with the view of Garry Kasparov (the first World Champion to lose a chess game against a computer) in advocating for a friendly collaboration between humans and AI (De Cremer & Kasparov, 2021). However, this paper aims to offer an additional insight. Not only will multiple artificial intelligences arise to solve different tasks, but also multiple artificial intelligences incorporating different combinations of human competences will emerge to execute the same task in different ways. This is especially true in scientific endeavors directed towards understanding the abstract cause of complex biological phenomena. Multiple organizational recombinations will emerge to capture different nuances of the infinite complexity of reality.

Past literature has shown that companies can hinder subsequent innovation (Williams, 2013), find it convenient to incorporate R&D that originates outside the firm (Chesbrough, 2003), are shifting away from investments in scientific research (Arora, Belenzon & Patacconi, 2018), and tend to focus on applied research with immediate commercial applications, while universities specialize

in basic research (Nelson, 1959). The AlphaFold case shows instead that companies can participate actively in the production of basic science without immediate commercial applications. Academics participating in the CASP initiative in 2018 had the sense that DeepMind was “in it for the science” (AlQuraishi, 2018, “AlphaFold” section, para. 1). A question necessarily arises: why would a company have the incentive to participate in a scientific competition such as the CASP initiative? Google’s DeepMind was not the only company participating in the CASP initiative. Other companies have participated in the CASP initiative, including Microsoft and Tencent. This question can have multiple answers. These companies may decide to participate in these scientific initiatives to fulfill their corporate social responsibility goals (Kaplan, 2019). Or they may be implementing novel non-market strategies (Dorobantu, Kaul & Zelner, 2017), especially in a period in which Big Tech is suffering from a severe loss of legitimacy. These companies may also be trying to validate their core competences (Prahalad, 1993), while spreading the cost of R&D efforts for producing this basic research over multiple applications (Cohen & Klepper, 1996). These factors may explain why Google might have decided to join the CASP initiative, “solve” the “protein folding” problem, and participate in a collective effort of scientific knowledge production.

The AlphaFold1 case shows that multiple stakeholders are needed for solving the Grand Challenges of the 21st century (McGahan, 2015). The scientific community may provide the foundational insights for understanding a complex biological phenomenon such as “protein folding” (Anfinsen, 1973), and all the theoretical contributions upon which DeepMind relied. Companies like DeepMind may instead provide the necessary resources and capabilities for processing huge amounts of data. At the end of the 13th edition of the CASP initiative in 2018, academics belittled “DeepMind’s contribution by noting its seeming incrementality and crediting their success to Alphabet’s resources” (AlQuraishi, 2018, “What just happened?” section, para. 2). The AlphaFold1 case shows that multiple actors – with diverse perspectives, goals, resources, and capabilities – can coalesce (Cavalli &

McGahan, 2023 forthcoming) for achieving one of the greatest goals posed to humanity: understanding the origins of life. To capitalize on the advances brought by AlphaFold1, and to produce better knowledge for achieving these goals, the principal investigators of frontier academic labs adapted by making their own labs larger and engaging in different knowledge-generating strategies based on their academic background. Artificial Intelligence tools such as AlphaFold will allow scientists not only to expand, but also to realize the increasing complexity of frontier science.

REFERENCES

- Abriata, L. 2021, July 21. *Google colab notebooks are already running Deepmind's AlphaFold v. 2*. Towards Data Science. <https://towardsdatascience.com/google-colab-notebooks-are-already-running-deepminds-alphafold-v-2-92b4531ec127>
- Acemoglu, D., & Restrepo, P. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2): 3-30.
- Agrawal, A., Gans, J., & Goldfarb, A. 2018. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Review Press.
- Agrawal, A., Gans, J., & Goldfarb, A. 2022. *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.
- Agrawal, A., Goldfarb, A., & Teodoridis, F. 2016. Understanding the changing structure of scientific inquiry. *American Economic Journal: Applied Economics*, 8(1): 100-128.
- Agrawal, A., McHale, J., & Oettl, A. 2018. Finding needles in haystacks: Artificial intelligence and recombinant growth. In *The economics of artificial intelligence: An agenda* (pp. 149-174). University of Chicago Press.
- Agrawal, A., McHale, J., & Oettl, A. 2022. *Superhuman science: How artificial intelligence may impact innovation*.
- AlQuraishi, M. 2018, December 9. *AlphaFold @ CASP13: 'What just happened?'*. <https://moalquraishi.wordpress.com/2018/12/09/alphafold-casp13-what-just-happened/>
- AlQuraishi, M. 2019. AlphaFold at CASP13. *Bioinformatics*, 35(22): 4862-4865.
- Anfinsen, C. B. 1973. Principles that govern the folding of protein chains. *Science*, 181(4096): 223-230.
- Arora, A., Belenzon, S., & Pataconi, A. 2018. The decline of science in corporate R&D. *Strategic Management Journal*, 39(1): 3-32.
- Asgari, N., Singh, K., & Mitchell, W. 2017. Alliance portfolio reconfiguration following a technological discontinuity. *Strategic Management Journal*, 38(5): 1062-1081.
- Athey, S. C., Bryan, K. A., & Gans, J. S. 2020, May. The allocation of decision authority to human and artificial intelligence. In *AEA Papers and Proceedings*, Vol. 110: 80-84.
- Azoulay, P. 2019. Small-team science is beautiful. *Nature*, 566(7744): 330-332.
- Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. 2020. Are ideas getting harder to find?. *American Economic Review*, 110(4): 1104-44.
- Bloom, N., Schankerman, M., & Van Reenen, J. 2013. Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4): 1347-1393.
- Boudreau, K.J., Guinan, E.C., Lakhani, K.R. & Riedl, C. 2016. Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Management Science*, 62(10), pp.2765-2783.
- Boudreau, K.J., Lacetera, N. & Lakhani, K.R. 2011. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science*, 57(5), pp.843-863.
- Boudreau, K.J., Lakhani, K.R. & Menietti, M. 2016. Performance responses to competition across skill levels in rank-order tournaments: field evidence and implications for tournament design. *The RAND Journal of Economics*, 47(1), pp.140-165.
- Bresnahan, T. 2023. What innovation paths for AI to become a GPT? *Journal of Economics & Management Strategy*.
- Bresnahan, T. F., & Trajtenberg, M. 1995. General purpose technologies 'Engines of growth?'. *Journal of econometrics*, 65(1): 83-108.
- Bryan, K.A. 2017. The perils of path dependence. *Survive and Thrive: Winning Against Strategic Threats to Your Business*, 111.
- Callaway, E. 2020, November 30. 'It will change everything': DeepMind's AI makes gigantic leap in solving protein structures. *Nature*, 588(7837): 203-205. <https://www.nature.com/articles/d41586-020-03348-4>
- Callaway, E. 2021, July 15. DeepMind's AI for protein structure is coming to the masses. *Nature*. <https://www.nature.com/articles/d41586-021-01968-y>
- Callaway, E. 2022, December 13. After AlphaFold: protein-folding contest seeks next big breakthrough. *Nature*. <https://www.nature.com/articles/d41586-022-04438-1>
- Castelvecchi, D. 2016, October 5. Can we open the black box of AI?. *Nature News*, 538(7623): 20. <https://www.nature.com/news/can-we-open-the-black-box-of-ai-1.20731>
- Cavalli, G. & McGahan, A. M. forthcoming 2023. Opening Innovation to Address Grand Challenges in Henry Chesbrough, Agnieszka Radziwon, Wim Vanhaverbeke, Joel West, eds. *Open Innovation Handbook* (Oxford University Press).

- Chesbrough, H. W. 2003. *Open innovation: The new imperative for creating and profiting from technology*. Harvard Business Press.
- Cockburn, I. M., Henderson, R., & Stern, S. 2019. 4. The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis. In *The economics of artificial intelligence*, 115-148. University of Chicago Press.
- Cohen, W. M., & Klepper, S. 1996. Firm size and the nature of innovation within industries: the case of process and product R&D. *The Review of Economics and Statistics*, 232-243.
- De Cremer, D., & Kasparov, G. 2021. AI should augment human intelligence, not replace it. *Harvard Business Review*, 18.
- DeepMind. 2020a, January 15. *AlphaFold: Using AI for scientific discovery*. <https://www.deepmind.com/blog/alphafold-using-ai-for-scientific-discovery-2020>
- DeepMind. 2020b, November 30. *AlphaFold: The making of a scientific breakthrough* [Video]. YouTube. <https://www.youtube.com/watch?v=gg7WjuFs8F4>
- DeepMind. 2020c, November 30. *AlphaFold: a solution to a 50-year-old grand challenge in biology*. <https://www.deepmind.com/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology>
- DeepMind. 2020d, November 30. *Developing an AI solution to 50-year-old protein challenge*. Phys.org. <https://phys.org/news/2020-11-ai-solution-year-old-protein.html>
- Dorobantu, S., Kaul, A., & Zelner, B. 2017. Nonmarket strategy research through the lens of new institutional economics: An integrative review and future directions. *Strategic Management Journal*, 38(1): 114-140.
- Fjelland, R. 2020. Why general artificial intelligence will not be realized. *Humanities and Social Sciences Communications*, 7(1), pp.1-9.
- Fleming, L. 2001. Recombinant uncertainty in technological search. *Management Science*, 47(1): 117-132.
- Foldit. 2020, January 31. *AlphaFold: Machine learning for protein structure prediction*. <https://fold.it/forum/blog/alphafold-machine-learning-for-protein-structure-prediction>
- Galasso, A., & Luo, H. 2018. Punishing Robots: issues in the economics of tort liability and innovation in artificial intelligence. In *The economics of artificial intelligence: An agenda*, 493-504. University of Chicago Press.
- Gambardella, A., & McGahan, A. M. 2010. Business-model innovation: General purpose technologies and their implications for industry structure. *Long Range Planning*, 43(2-3): 262-271.
- Griliches, Z. 1957. Hybrid corn: An exploration in the economics of technological change. *Econometrica, Journal of the Econometric Society*, 501-522.
- Gross, D.P. 2020. Creativity under fire: The effects of competition on creative production. *Review of Economics and Statistics*, 102(3), pp.583-599.
- Haeussler, C. & Sauermann, H. 2020. Division of labor in collaborative knowledge production: The role of team size and interdisciplinarity. *Research Policy*, 49(6), p.103987.
- Jeppesen, L. B., & Lakhani, K. R. 2010. Marginality and problem-solving effectiveness in broadcast search. *Organization Science*, 21(5): 1016-1033.
- Jones, B. F. 2009. The burden of knowledge and the “death of the renaissance man”: Is innovation getting harder?. *The Review of Economic Studies*, 76(1): 283-317.
- Jones, B. F. 2021. The rise of research teams: Benefits and costs in economics. *Journal of Economic Perspectives*, 35(2): 191-216.
- Jones, D. T., & Thornton, J. M. 2022, January 11. The impact of AlphaFold2 one year on. *Nature Methods*, 19(1): 15-20. <https://www.nature.com/articles/s41592-021-01365-3>
- Kaplan, S. 2019. The 360° Corporation. In *The 360° Corporation*. Stanford University Press.
- Kaplan, S., & Tripsas, M. 2008. Thinking about technology: Applying a cognitive lens to technical change. *Research Policy*, 37(5): 790-805.
- Kaplan, S., & Vakili, K. 2015. The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal*, 36(10): 1435-1457.
- Kim, S. 2022. Shortcuts to Innovation: The Use of Analogies in Knowledge Production. In *Academy of Management Proceedings*, Vol. 2022, No. 1: 10675. Briarcliff Manor, NY 10510: Academy of Management.
- King, A.A. & Lakhani, K.R. 2011. The contingent effect of absorptive capacity: an open innovation analysis. *Boston: Harvard Business School*.
- Kitcher, P. 1990. The division of cognitive labor. *The Journal of Philosophy*, 87(1): 5-22.
- Korinek, A., & Stiglitz, J. E. 2018. Artificial intelligence and its implications for income distribution and unemployment. In *The Economics of Artificial Intelligence: An Agenda*, 349-390. University of Chicago Press.
- Kurzweil, R. 2005. *The singularity is near: When humans transcend biology*. Penguin.

- Lacetera, N., Cockburn, I. M., & Henderson, R. 2004. Do firms change capabilities by hiring new people? A study of the adoption of science-based drug discovery. In *Business strategy over the industry lifecycle*. Emerald Group Publishing Limited.
- Lane, J.N., Ganguli, I., Gaule, P., Guinan, E. & Lakhani, K.R. 2021. Engineering serendipity: When does knowledge sharing lead to knowledge production? *Strategic Management Journal*, 42(6), pp.1215-1244.
- Latour, B., & Woolgar, S. 2013. *Laboratory life: The construction of scientific facts*. Princeton University Press.
- Leahey, E., Beckman, C. M., & Stanko, T. L. 2017. Prominent but less productive: The impact of interdisciplinarity on scientists' research. *Administrative Science Quarterly*, 62(1): 105-139.
- LeCun, Y., Bengio, Y., & Hinton, G. 2015. Deep learning. *Nature*, 521(7553): 436-444.
- Leiponen, A., & Helfat, C. E. 2010. Innovation objectives, knowledge sources, and the benefits of breadth. *Strategic Management Journal*, 31(2): 224-236.
- Levinthal, D. A. 1997. Adaptation on rugged landscapes. *Management Science*, 43(7): 934-950.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organization Science*, 2(1): 71-87.
- McGahan, A. M. 2015. The Power of Coordination Action: Short-term Organizations with Long-term Impact. *Rotman Magazine*, 7-11.
- Merton, R. K. 1968. The Matthew Effect in Science: The reward and communication systems of science are considered. *Science*, 159(3810): 56-63.
- Moore, P. B., Hendrickson, W. A., Henderson, R., & Brunger, A. T. 2022. The protein-folding problem: Not yet solved. *Science*, 375(6580): 507-507. <https://www.science.org/doi/10.1126/science.abn9422>
- Moser, P. & Voena, A. 2012. Compulsory licensing: Evidence from the trading with the enemy act. *American Economic Review*, 102(1), pp.396-427.
- Nelson, R. R. 1959. The simple economics of basic scientific research. *Journal of Political Economy*, 67(3): 297-306.
- Orlikowski, W. J. 1992. The duality of technology: Rethinking the concept of technology in organizations. *Organization Science*, 3(3): 398-427.
- Orlikowski, W. J. 2000. Using technology and constituting structures: A practice lens for studying technology in organizations. *Organization Science*, 11(4): 404-428.
- Pammolli, F., Magazzini, L., & Riccaboni, M. 2011. The productivity crisis in pharmaceutical R&D. *Nature Reviews Drug Discovery*, 10(6): 428-438.
- Perignat, E., & Fleming, F. F. 2022. Sunk-Cost Bias and Knowing When to Terminate a Research Project. *Angewandte Chemie*, 134(36): e202208429.
- Prahalad, C. K. 1993. The role of core competencies in the corporation. *Research-Technology Management*, 36(6): 40-47.
- Puranam, P. 2021. Human-AI collaborative decision-making as an organization design problem. *Journal of Organization Design*, 10(2), pp.75-80.
- Sampat, B. N., & Lichtenberg, F. R. 2011. What are the respective roles of the public and private sectors in pharmaceutical innovation?. *Health Affairs*, 30(2): 332-339.
- Schilling, M. A., & Green, E. 2011. Recombinant search and breakthrough idea generation: An analysis of high impact papers in the social sciences. *Research Policy*, 40(10): 1321-1331.
- Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., ... & Hassabis, D. 2020. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792): 706-710.
- Silva, J. S., & Tenreiro, S. 2006. The log of gravity. *The Review of Economics and Statistics*, 88(4): 641-658.
- Stiglitz, J. E. 1999. Knowledge as a global public good. *Global public goods: International cooperation in the 21st century*, 308: 308-325.
- Stiglitz, J. E., McFadden, D., & Peltzman, S. 1987. Technological change, sunk costs, and competition. *Brookings Papers on Economic Activity*, 1987(3): 883-947.
- Strevens, M. 2003. The role of the priority rule in science. *The Journal of Philosophy*, 100(2): 55-79.
- Taylor, A., & Greve, H. R. 2006. Superman or the fantastic four? Knowledge combination and experience in innovative teams. *Academy of Management Journal*, 49(4): 723-740.
- Teodoridis, F. 2018. Understanding team knowledge production: The interrelated roles of technology and expertise. *Management Science*, 64(8): 3625-3648.
- Teodoridis, F., Bikard, M., & Vakili, K. 2019. Creativity at the knowledge frontier: The impact of specialization in fast-and slow-paced domains. *Administrative Science Quarterly*, 64(4): 894-927.

- Toews, R. 2021, October 3. *AlphaFold Is The Most Important Achievement In AI—Ever*. Forbes.
<https://www.forbes.com/sites/robtoews/2021/10/03/alphafold-is-the-most-important-achievement-in-ai-ever/?sh=2bc94f626e0a>
- Turing, A. M. 1950. Computing Machinery and Intelligence. *Mind*, 59(236): 433–460.
- Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. 2013. Atypical combinations and scientific impact. *Science*, 342(6157): 468-472.
- Vakili, K., & Kaplan, S. 2021. Organizing for innovation: A contingency view on innovative team configuration. *Strategic Management Journal*, 42(6): 1159-1183.
- Weisberg, R. W. 1999. Creativity and knowledge: A challenge to theories. In R. J. Sternberg (Ed.), *Handbook of Creativity*, 226–250. Cambridge University Press.
- Weitzman, M. L. 1998. Recombinant growth. *The Quarterly Journal of Economics*, 113(2): 331-360.
- Williams, H. L. 2013. Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy*, 121(1): 1-27.
- Wuchty, S., Jones, B. F., & Uzzi, B. 2007. The increasing dominance of teams in production of knowledge. *Science*, 316(5827): 1036-1039.
- Zhang, Y., & Yang, W. 2022. Breakthrough Invention and Problem Complexity: Evidence from a Quasi-Experiment. *Strategic Management Journal*.
- Zhuo, R., 2022. Exploit or Explore? An Empirical Study of Resource Allocation in Scientific Labs. *Job Market Paper*.

FIGURES AND TABLES

Figure 1: Decision Tree for Principal Investigator of Lab k

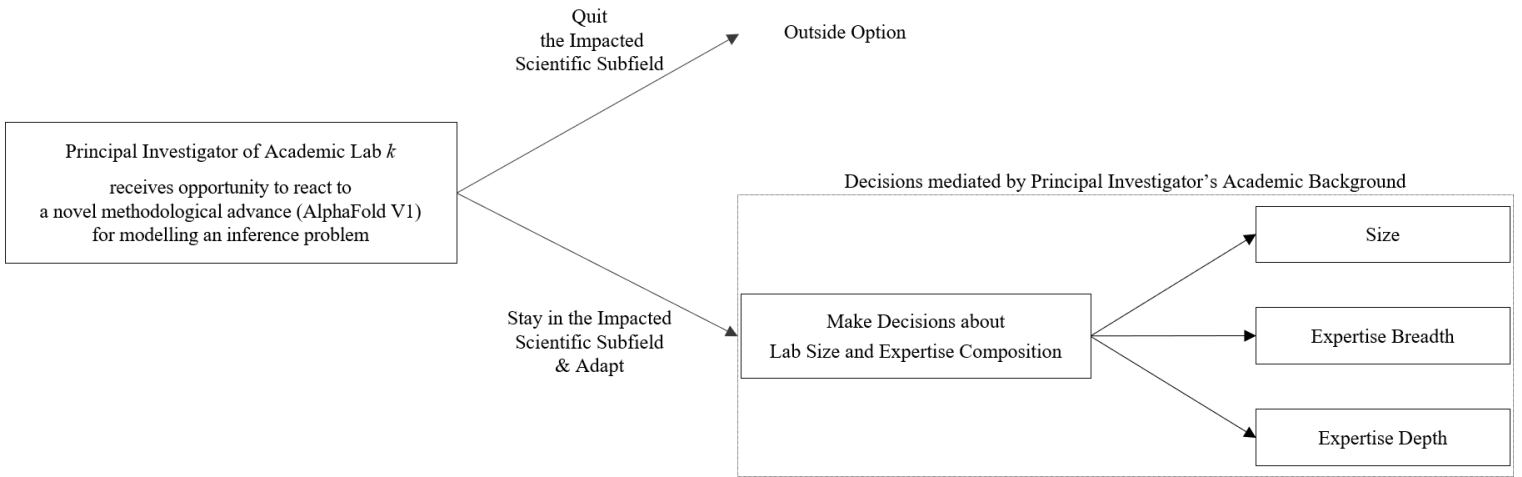


Figure 2: Locations of the academic institutions in which established academic labs are based (177 Labs; 120 academic institutions [blue dots]; 25 countries)



Figure 3: Established academic labs by subfield and country

#	Country	Labs engaged in PF-related software development	Labs engaged in software development for other protein-related phenomena	Total Labs per Country
1	Argentina	0	1	1
2	Belgium	0	1	1
3	Brazil	1	0	1
4	Canada	0	1	1
5	China	2	2	4
6	France	3	3	6
7	Germany	2	10	12
8	Hungary	0	1	1
9	India	1	1	2
10	Ireland	1	0	1
11	Israel	1	7	8
12	Italy	0	2	2
13	Japan	1	0	1
14	Lithuania	1	0	1
15	Netherlands	2	2	4
16	Poland	1	2	3
17	Portugal	0	1	1
18	Saudi Arabia	0	2	2
19	South Korea	1	0	1
20	Spain	3	5	8
21	Sweden	2	0	2
22	Switzerland	1	2	3
23	Turkey	0	1	1
24	United Kingdom	4	12	16
25	United States	28	66	94
Total		55	122	177

Figure 4: Institutions by country and labs' scientific subfield


#	Country	Institutions with labs engaged in PF-related software development	Institutions with labs engaged in both PF-related and other protein-related phenomena	Institutions with labs engaged in software development for other protein-related phenomena	Total institutions per country
1	Argentina			National University of Quilmes	1
2	Belgium			KU Leuven	1
3	Brazil	Laboratorio Nacional de Computacao Cientifica			1
4	Canada			University of Toronto	1
5	China	Nankai University	Huazhong University of Science and Technology	Westlake University	3
6	France	Institut Universitaire de France; Laboratoire Jean Kuntzmann	University of Paris	Conservatoire National des Arts et Métiers; Institut de Biologie Paris-Seine	5
7	Germany	Heinrich Heine Universität Düsseldorf; Technical University of Berlin		European Molecular Biology Laboratory (Hamburg); Forschungszentrum Jülich; Martin Luther University Halle-Wittenberg; Max Planck Institute for Developmental Biology; Max Planck Institute of Biophysics; Technical University of Munich; Universität Hamburg; University of Bayreuth	10
8	Hungary			Eötvös Loránd University	1
9	India	Indraprastha Institute of Information Technology Delhi		Indian Institute of Science Education and Research Pune	2
10	Ireland	University College Dublin			1
11	Israel	Ben-Gurion University of the Negev		Hebrew University of Jerusalem; Tel Aviv University; Weizmann Institute of Science	4
12	Italy			University of Bologna; University of Milan	2
13	Japan	RIKEN			1
14	Lithuania	Vilnius University			1
15	Netherlands	Radboud University Nijmegen; Utrecht University		Antoni van Leeuwenhoek Hospital; VU Amsterdam	4
16	Poland		University of Gdansk	University of Warsaw	2
17	Portugal			University of Coimbra	1
18	Saudi Arabia			King Abdullah University of Science and Technology	1
19	South Korea	Seoul National University			1
20	Spain	Barcelona Supercomputing Center;		Autonomous University of Barcelona; Centre for Genomic Regulation; University of Barcelona	5

		Pompeu Fabra University			
21	Sweden	Linköping University; Stockholm University			2
22	Switzerland	University of Basel		École Polytechnique Fédérale de Lausanne	2
23	Turkey			Middle East Technical University	1
24	United Kingdom	The Francis Crick Institute; University of Oxford; University of Reading	University College London	European Bioinformatics Institute; Imperial College London; Medical Research Council; University of Cambridge; University of London; University of Manchester	10
25	United States	Brigham Young University; City of Hope; Michigan State University; Nationwide Children's Hospital; Pacific Lutheran University; Rice University; Stony Brook University; University at Buffalo, State University of New York; University of Florida; University of Miami; University of Michigan-Ann Arbor; University of Missouri; University of New Orleans; Western Washington University	Columbia University; Florida State University; Georgia Institute of Technology; Purdue University West Lafayette; Rutgers, The State University of New Jersey; University of Chicago; University of Kansas; University of Washington; Virginia Tech	Albert Einstein College of Medicine; Baylor College of Medicine; Boston University; Clemson University; Colorado State University; Cornell University; Dartmouth College; Fox Chase Cancer Center; George Mason University; Harvard University; Indiana University; Iowa State University; Johns Hopkins University; Louisiana State University; Massachusetts Institute of Technology; Old Dominion University; Stanford University; Texas A&M University; The University of Texas at Austin; The University of Texas Southwestern Medical Center; University of California, Berkeley; University of California, Irvine; University of California, Los Angeles; University of California, Riverside; University of California, San Francisco; University of Delaware; University of Georgia; University of Illinois at Chicago; University of Illinois Urbana-Champaign; University of Maryland, College Park; University of Massachusetts Amherst; University of North Carolina at Chapel Hill; University of Pittsburgh; Worcester Polytechnic Institute	57
Total		36	13	71	120

Figure 5: Treated and control groups, for all labs and senior labs, by principal investigators' academic background

	All labs			Senior labs (operational 2015-2021)		
	Treated	Control	Total	Treated	Control	Total
PI LS	13	32	45	12	25	37
PI CS	14	23	37	9	21	30
PI LS+CS	28	67	95	21	55	76
Total	55	122	177	42	101	143

Figure 6: Established academic labs engaged in PF-related software development (i.e., labs that have participated in CASP at least once in the period 2014-2018) increased their average size through recruitment of PhD students and Postdoctoral researchers following the unexpected success of AlphaFold1 in 2018.



Figures 7a, 7b, 7c, 7d (All labs): Yearly difference-in-difference coefficients show increases in size in terms of PhD students and Postdoctoral researchers within established academic labs engaged in PF-related software development, following the unexpected success of AlphaFold1 in 2018 (Figure 7b).

Figure 7a:




Figure 7b:

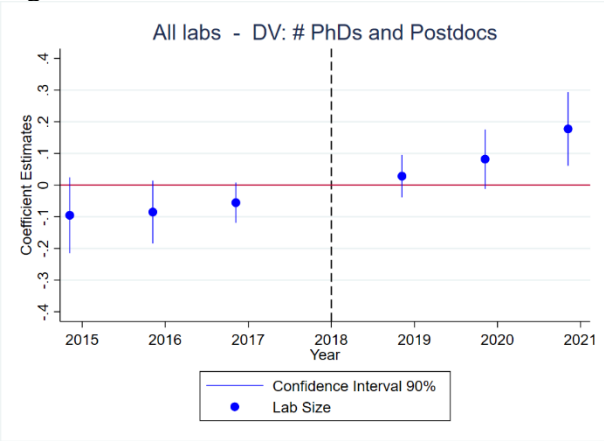


Figure 7c:





Figure 7d:



Figures 8a, 8b, 8c, 8d (Senior labs): Yearly difference-in-difference coefficients show increases in size in terms of PhD students and Postdoctoral researchers within established academic labs engaged in PF-related software development, following the unexpected success of AlphaFold1 in 2018 (Figure 8b)

Figure 8a:




Figure 8b:




Figure 8c:


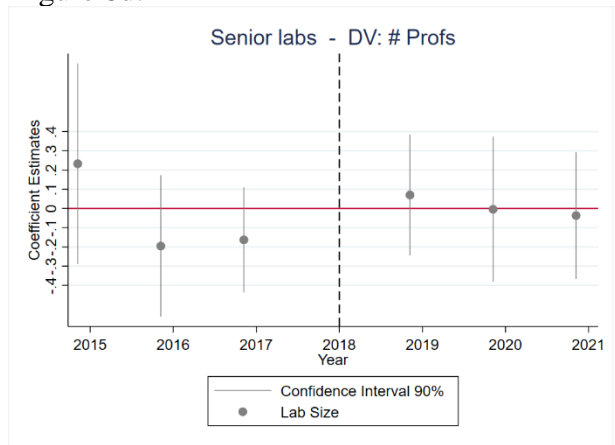


Figure 8d:



Figures 9a, 9b, 9c (All labs; focus on PhD students and Postdoctoral researchers): Yearly triple difference coefficients show increases in size for PF labs led by principal investigators with a CS academic background (Figure 9a) and decreases in size for PF labs led by principal investigators with a LS+CS academic background (Figure 9c), following the unexpected success of AlphaFold1 in 2018.

Figure 9a:




Figure 9b:





Figure 9c:



Figures 10a, 10b, 10c (Senior labs; focus on PhD students and Postdoctoral researchers): Yearly triple difference coefficients show increases in size for PF labs led by principal investigators with a CS background (Figure 10a) and decreases in size for labs led by principal investigators with a LS+CS background (Figure 10c), following the unexpected success of AlphaFold1 in 2018.

Figure 10a:




Figure 10b:





Figure 10c:



Figures 11a, 11b, 11c (All labs; focus on CS Principal Investigators): Principal investigators of PF labs with a CS academic background increased recruitment of lab members with a CS background (Figure 11a), and stopped recruitment of lab members with an LS background (Figure 11b), following the unexpected success of AlphaFold1 in 2018.

Figure 11a:




Figure 11b:

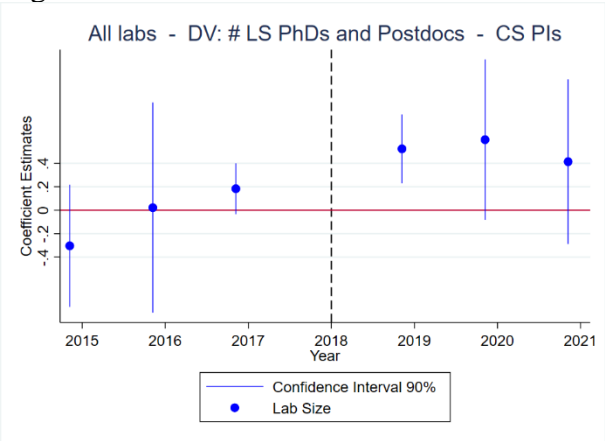



Figure 11c:



Figures 12a, 12b, 12c (All labs; focus on LS Principal Investigators): Principal investigators of PF labs with an LS academic background increased recruitment of lab members with a CS background (Figure 12a), and stopped recruitment of lab members with an LS background (Figure 12b), following the unexpected success of AlphaFold1 in 2018

Figure 12a:




Figure 12b:





Figure 12c:



Figures 13a, 13b, 13c (All labs; focus on LS+CS Principal Investigators): Principal investigators of PF labs with an LS+CS academic background reduced recruitment of lab members with a CS background (Figure 13a), and stopped dismissing lab members with an LS background (Figure 13b), following the unexpected success of AlphaFold1 in 2018.

Figure 13a:




Figure 13b:




Figure 13c:




Table 1: Principal investigators of established academic labs engaged in PF-related software development (i.e., labs that have participated in CASP at least once in the period 2014-2018) have adapted by recruiting more PhD students and Postdoctoral researchers following the unexpected success of AlphaFold1 in 2018 (Columns 5 and 6; result in Column 6 is significant at $p < 0.111$).

DV: Lab Size
Poisson Pseudo-Maximum Likelihood Estimates, Period: 2015-2021

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	# All Members		# Undergrads & Masters		# PhDs & Postdocs		# Staff		# Professors	
	All	Senior	All	Senior	All	Senior	All	Senior	All	Senior
	labs	labs	labs	labs	labs	labs	labs	labs	labs	labs
PostAlphaFoldV1 _t · Lab _k PF _k	0.053 (0.056)	0.033 (0.059)	-0.117 (0.129)	-0.109 (0.141)	0.153** (0.069)	0.113 (0.071)	-0.069 (0.111)	-0.067 (0.117)	0.059 (0.159)	0.036 (0.164)
Observations	1,190	1,001	936	798	1,167	987	629	567	374	336
Observed labs	177	143	139	114	173	141	92	81	55	48

Robust standard errors in parentheses. Standard errors are clustered at the lab level.

All columns include lab and year fixed effects.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$