

DECEMBER 2020

Mapping U.S. Multinationals' Global AI R&D Activity

CSET Issue Brief



AUTHORS

Roxanne Heston
Remco Zwetsloot

Table of Contents

Executive Summary	2
1. Introduction: U.S. Tech Multinationals and AI	4
2. Methodology and Data	5
3. Findings.....	8
Where Are AI Labs and Staff Located?.....	8
What Types of R&D Do AI Labs Conduct?	14
What Drives Location Decisions?.....	19
Access to Talent	19
Market Access and Product Adaptation.....	21
Cost Savings.....	22
Conclusion	24
Acknowledgments	26
Appendix: Data and Methods	27
A. Locating AI Labs	27
B. Locating AI Staff.....	28
C. Comparing AI Labs, AI Staff, and Other R&D Measures	30
D. Classifying AI R&D Activities	32
Endnotes.....	34

Executive Summary

This paper presents new data on the global distribution of U.S. tech companies' artificial intelligence labs and staff. It focuses on six companies—Amazon, Apple, Facebook, Google, IBM, and Microsoft—all of which have a history of conducting cutting-edge AI research and development. Our findings shed light on where these companies conduct AI R&D and why they select particular locations.

This paper also addresses debates about the benefits and risks to national security of having U.S. companies conduct R&D abroad. On its own, new data will not solve these debates, but data is a prerequisite for nuanced policy discussions and cost-benefit analyses. To facilitate future research on the topic, we are releasing our new global AI labs dataset [on CSET's GitHub page](#).

Key findings include:

- Company AI labs are spread all over the world, especially in North America, Europe, and Asia.
 - Media coverage frequently mentions “AI labs,” but not all companies conduct R&D along a “lab” structure. Amazon, for example, integrates its R&D with the work of its product teams and has no standalone labs. Apple, alternatively, does not release information about its R&D activities.
 - For the four companies where we could find information on labs—Facebook, Google, IBM, and Microsoft—we found 62 labs conducting AI R&D. The majority of these labs (68 percent) were located outside of the United States.
 - European countries, mainly the United Kingdom and France, host 19 percent of AI labs in our dataset. Other common destinations were Israel and China (10 percent each) and India (8 percent). The only region in the world without any AI labs was Latin America.
- In contrast to AI labs, most company AI staff remain concentrated in the United States.

- Whereas most AI labs are located abroad, 68 percent of AI staff at these companies are located within the United States. The San Francisco Bay Area and the Seattle region—where the headquarters of five out of the six companies we focus on are located—host nearly half of these companies’ global AI staff.
- The companies combined have 500-plus AI staff in 30 locations across 15 countries. The main company hubs are almost exclusively in large cities, both at home and abroad.
- Companies mainly focus on AI research in North America and Europe, and on AI development in other regions.
 - About 85 percent of labs in North America and Europe do some form of basic AI research, and less than 15 percent focus on development. In contrast, a large majority of labs in other countries, such as India and Israel, focus on applied research or development.
 - Some companies’ labs appear to focus on research, whereas others focus on development. Most labs conduct multiple types of R&D.
- Companies decide to conduct AI R&D abroad for multiple reasons, but access to global talent appears to be the dominant factor.
 - According to company statements and industry analysts, the main reason that companies set up AI R&D labs abroad is access to talent. Other reasons include cost savings, as well as market access and product adaptation.
 - Media reports and economic studies suggest U.S. immigration restrictions are associated with U.S. companies’ global expansion. Our findings similarly suggest that immigration restrictions push companies to set up AI labs abroad. At the same time, overseas expansion is unlikely to stop entirely even if U.S. immigration policy became less restrictive.

1. Introduction: U.S. Tech Multinationals and AI

The private sector is a driving force behind technical progress in AI and machine learning, and a large fraction of private sector AI R&D is done by a limited number of large tech companies. A McKinsey report that surveyed 35 major technology companies found that they spent \$18-27 billion in internal funds on AI development in 2016, and another \$8-12 billion on investments and acquisitions.¹ Since 2010, big tech in the United States absorbed more than 60 smaller AI companies.² The President's Council of Advisors on Science and Technology predicts that U.S. companies will spend more than \$100 billion on AI R&D annually by 2025.³ By comparison, the U.S. federal government intends to spend about \$5 billion on unclassified AI and machine learning-related research and development in FY2020.⁴

Many of the world's biggest AI companies are American; this paper focuses on Amazon, Apple, Facebook, Google, IBM, and Microsoft in particular. Together, these companies spend over \$76 billion on R&D annually, and their collective market capitalization stands above \$5 trillion.⁵ Their business models are globalized. Apple, Facebook, Google, and IBM each receive less than half of their total revenue from the U.S. market. Microsoft and Amazon receive 51 percent and 69 percent, respectively.⁶

The global reach of U.S. tech companies provides several benefits to the United States. The fact that they operate across the world allows them to innovate at a scale and speed that would be impossible if they were focused solely on domestic markets, and ensures they can absorb talent and ideas from all over the world. The U.S. economy and national security both benefit. Some economists argue that U.S. multinationals expanding abroad contribute to more domestic employment, exports, and R&D, and that "less investment abroad by US firms would weaken—not strengthen—the US economy."⁷

But these companies' scale and reach also come with risks. From the perspective of national security, for example, the commercial sector's global presence is a double-edged sword. As the 2018 National Defense Strategy put it, "the fact that many technological developments will come from the commercial sector means that state competitors and non-state actors will also have access to them, a fact that risks eroding the conventional overmatch to which our nation has grown accustomed."⁸ U.S. multinationals' R&D activities abroad can boost other countries' technology ecosystems in several ways, for example, by training local talent, or through joint ventures involving

intellectual property or collaborative research.⁹ One industry expert argues “if Microsoft had never founded [Microsoft Research Asia], that would have significantly delayed the rise of China’s AI ecosystem.”¹⁰ In 2019, the then-chairman of the Joint Chiefs of Staff said that Google’s work in China was “indirectly benefiting the Chinese military.”¹¹ Others, including Microsoft and Google, dispute such claims.¹²

These different benefits and concerns mean that U.S. policymakers face complex trade-offs as they think about whether and how to regulate U.S. companies’ AI activities abroad. Resolving these complex trade-offs requires data. To assess the implications of new policies or the continuation of the status quo, policymakers will need to know where U.S. multinationals are active, what activities they engage in abroad, and why. Analysts and decision-makers across the political spectrum have long argued that better data is needed on these and related questions.¹³

This paper informs these conversations by mapping the global AI R&D activities of Amazon, Apple, Facebook, Google, IBM, and Microsoft. Section 2 describes how we collected and analyzed new data on these companies’ global AI labs and AI staff. Section 3 then presents our findings on where company labs and staff are located, what types of R&D they appear to be engaged in across different locations, and what factors drive companies’ choice of location.

This research is meant to be a first step; many questions about U.S. companies’ AI R&D activities remain unanswered. To facilitate future research on the topic, we are releasing our new dataset of AI labs online along with the publication of this paper. The Appendix discusses methodological choices and other issues relevant to such future research.

2. Methodology and Data

This paper presents new data on the AI labs and staff of six big U.S. tech multinationals: Amazon, Apple, Facebook, Google, IBM, and Microsoft. These companies were selected based on their high score on several AI R&D-related metrics, including R&D spending, AI startup acquisitions, patent holdings, and involvement in professional AI organizations.¹⁴ For these companies, we sought to understand where they locate their AI R&D activities and why.

The first R&D metric we considered involved “labs,” which are featured prominently in media and industry analyses of companies’ AI R&D. We compiled lists of labs by searching companies’ own websites and news reports, and confirmed the lists’ accuracy with company representatives.* We found no public AI labs at Amazon and Apple; in 2014, Amazon started integrating its AI research staff into product-focused teams,¹⁵ while Apple does not publicly discuss its AI R&D activities.¹⁶ For the remaining four companies—Facebook, Google, IBM, and Microsoft—we found a total of 62 labs, of which a plurality (23) are run by Google (Figure 1).¹⁷ The list is current as of October 2019.

Figure 1. Number of public AI labs by company.



Source: CSET AI labs data (see Appendix A). Amazon and Apple are omitted because they do not use a “lab” structure or do not publicly discuss their labs.

The second metric we looked at is staff counts. We count “AI staff” as any company employee who listed one or more “AI skills,” such as “computer vision” (a category of current AI research) or “PyTorch” (a specific machine learning library), on their LinkedIn profile.† LinkedIn is, to our knowledge, the only data source on AI skills and jobs with global coverage. However, its data is not perfect; as discussed below, rates of LinkedIn usage vary across

* We adopted companies’ own definitions of “lab” and “artificial intelligence.” Different companies likely use these terms slightly differently; see the section of this paper on “What Types of R&D Do AI Labs Conduct,” as well as Appendix A, for more discussion of this issue.


† The full list of 46 AI skills we used to conduct our search is available in Appendix B.

countries, and data is available only as relative company totals and location-specific aggregates.¹⁸

Figure 2 compares the percentage of AI staff across the companies to the percentage of AI labs, and shows that the relative distribution of the two indicators differs considerably across companies. For example, while Google has more than a third of all labs, it has only 23 percent of all staff. Staff data are current as of May 2020.¹⁹

Figure 2. Global distribution of AI labs and staff by company.

SHARE OF TOTAL AT EACH COMPANY



Source: CSET AI labs data (see Appendix A); AI staff data from LinkedIn (see Appendix B).²⁰

Both the lab and staff metrics have strengths and weaknesses as measures of AI R&D activity. For example, the presence of an AI lab indicates that AI R&D is taking place, but the absence of a lab does not imply the absence of R&D. Snapshot measures of staff counts are a good indicator for current activity, but do not reveal whether a company plans to expand in a certain location or what they are working on. For instance, Facebook AI Research in Paris had just six researchers at its founding, but the company said from the start that it planned to expand the lab, which eventually grew more than ten-fold.²¹ Both measures sometimes suffer from lack of data availability, and results on one measure can help fill the gaps of the other. To provide as comprehensive a picture as possible, this report provides data on both labs and staff counts when such data is available. The Appendix describes our methodology in greater detail. The full CSET AI labs database is available at <https://github.com/georgetown-cset/us-multinational-ai-labs>.


3. Findings

This section presents our main findings. First, we discuss where Amazon, Apple, Facebook, Google, IBM, and Microsoft have located their AI R&D activities, analyzing data on both labs and staff. Second, for the four companies that have AI labs, we analyze what types of AI R&D they do in different locations. Third, we review what prior research and company statements suggest about the reasons that companies decide to locate AI R&D in particular places.

Where Are AI Labs and Staff Located?

In mapping U.S. multinational corporations' (MNC) AI R&D activity, we study two metrics: AI labs and AI staff. Figure 3 shows these two metrics are distributed differently—while only about one-third (32 percent) of AI labs are located in the United States, around two-thirds (68 percent) of AI staff are. Below, we will first take a closer look at the geographic distribution of labs, followed by more detail on the distribution of staff.

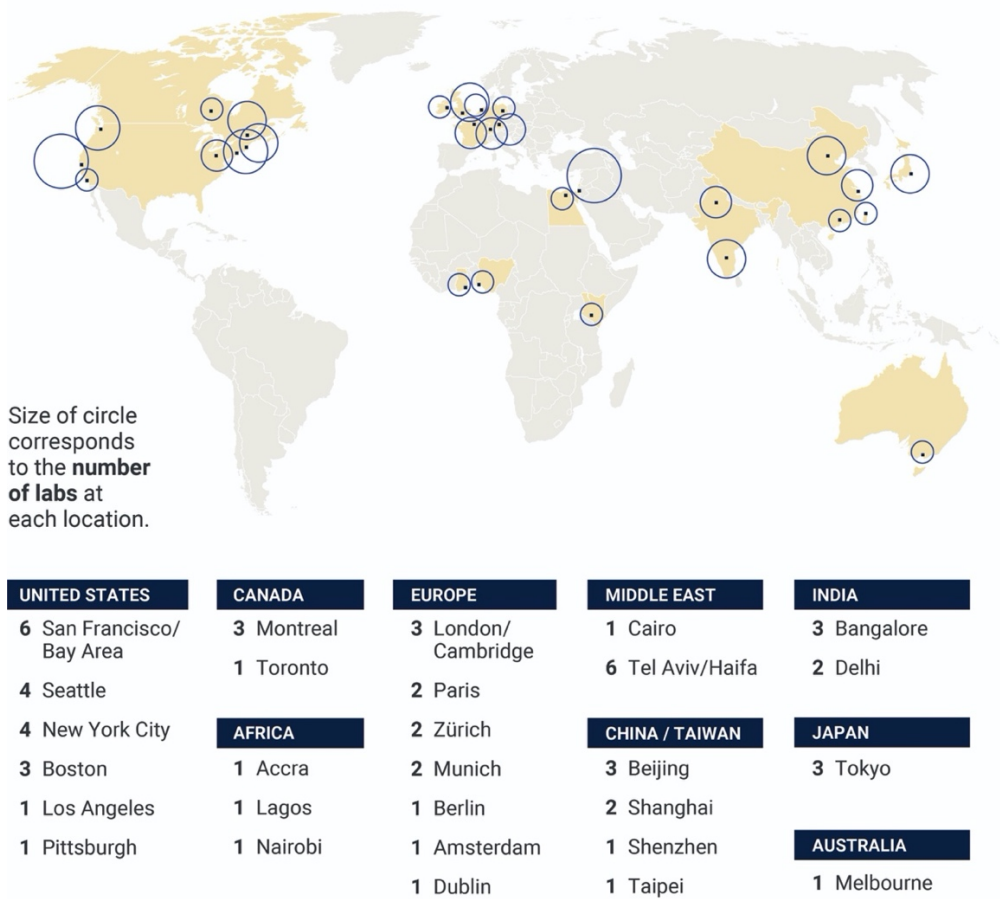
Figure 3. Percentage of U.S. MNC AI labs and staff by U.S. versus non-U.S. location.



Source: CSET AI labs data (see Appendix A); AI staff data from LinkedIn (see Appendix B). Included companies for staff are Amazon, Apple, Facebook, Google, IBM, and Microsoft; no labs data available for Amazon and Apple.

Figure 4 shows the international distribution of the 62 AI labs in our dataset, which are located in 17 countries and 41 cities across 29 metropolitan areas.²² Of the 20 labs located in the United States, six are located in Silicon Valley, four in or near New York City, and another four in or near Seattle. Outside of the United States, the most popular countries include Israel, China, and India. Outside of countries with established technological hotspots, we found four labs in Africa: one each in Egypt, Ghana, Kenya, and Nigeria. Notably, we did not find any labs located in Latin America.

Figure 4. Global distribution of AI labs by location.



Source: CSET AI labs data (see Appendix A).

Within countries, the labs are predominantly located in well-known “tech clusters,” such as London, Beijing, and the west coast region in Israel

commonly referred to as “Silicon Wadi.” In line with economics research on R&D clustering, the market areas that host the most labs are typically also home to active AI startup ecosystems and universities that have highly ranked AI programs.²³ Indeed, companies sometimes establish labs by acquiring those startups or partnering with universities (Box 1). Even if they do not, companies often benefit from being located there because clusters are home to large talent pools and new technical ideas (See section on “What Drives Location Decisions.”)²⁴

Box 1. How Companies Establish AI Labs

Labs in our dataset were established in one of three ways: by hiring or relocating staff internally, by partnering with local institutions, or by acquiring startups.

Internal hiring and relocation is by far the most common mode of establishment. Sometimes a company will create a standalone location specifically for work on AI. One such example is Google Brain Tokyo, one of more than a dozen Google Brain groups around the globe.²⁵ In other cases, the company will extend its existing (non-AI) R&D facilities to include AI-focused projects.

Companies can also partner with universities or other organizations to create joint R&D projects or facilities. One common approach within AI is for a company to partner with a professor at a prestigious university, as Facebook AI Research did with Carnegie Mellon University professors when it set up its Pittsburgh offices.²⁶ Institutional partnerships are also an option. The Massachusetts Institute of Technology (MIT) alone is home to two joint labs in our database: the MIT-IBM Watson AI Lab and Microsoft Research Lab - AI.²⁷ Meanwhile, in Shanghai, Microsoft teamed up with the state-owned enterprise INESA to create an AI Innovation Center, where it does regional AI product development.²⁸

Finally, companies will occasionally establish labs by acquiring startups. Since 2010, there have been at least 635 AI company acquisitions, roughly a tenth of which were made by the six companies that this report focuses on.²⁹ Microsoft’s acquisition of Montreal-based startup Maluuba, which formed the basis for a significant expansion of AI R&D in Montreal in 2018, is one example.³⁰ Another clear case of a startup turned lab is DeepMind. Since Google acquired the company in 2014, DeepMind has not only grown its headquarters to hundreds of staff but has also spawned five global offices. However, not every startup acquisition leads to a new lab. Startup staff and intellectual property are often folded into existing products or teams instead, in which case we do not count them as standalone “labs.”³¹

Figure 5 takes a closer look at the locations of AI labs across the four companies with labs. The United States is the most common location across all the companies, but each has labs across Europe and in Israel as well, and three of the four companies have labs in Canada, China, and India. In general, IBM and Microsoft appear proportionately more active outside of North America and Europe than Google and Facebook.*

Figure 5. Number of AI labs in different locations, by company.

	TOTAL LABS AT EACH LOCATION	SHARE OF ALL LABS	NUMBER OF LABS AT EACH LOCATION BY COMPANY			
			Facebook	Google	IBM	Microsoft
United States	20	32%	4 ■■■■	9 ■■■■ ■	3 ■■■	4 ■■■■
Europe	12	19%	2 ■■	6 ■■■■ ■	1 ■	3 ■■■
China	6	10%	0	1 ■	1 ■	4 ■■■■
Israel	6	10%	1 ■	2 ■■	1 ■	2 ■■
India	5	8%	0	1 ■	2 ■■	2 ■■
Canada	4	6%	1 ■	2 ■■	0	1 ■
Other	9	15%	0	2 ■■	3 ■■■	4 ■■■■
Total	62 labs		8 labs	23 labs	11 labs	20 labs

Source: CSET AI labs data (see Appendix A).

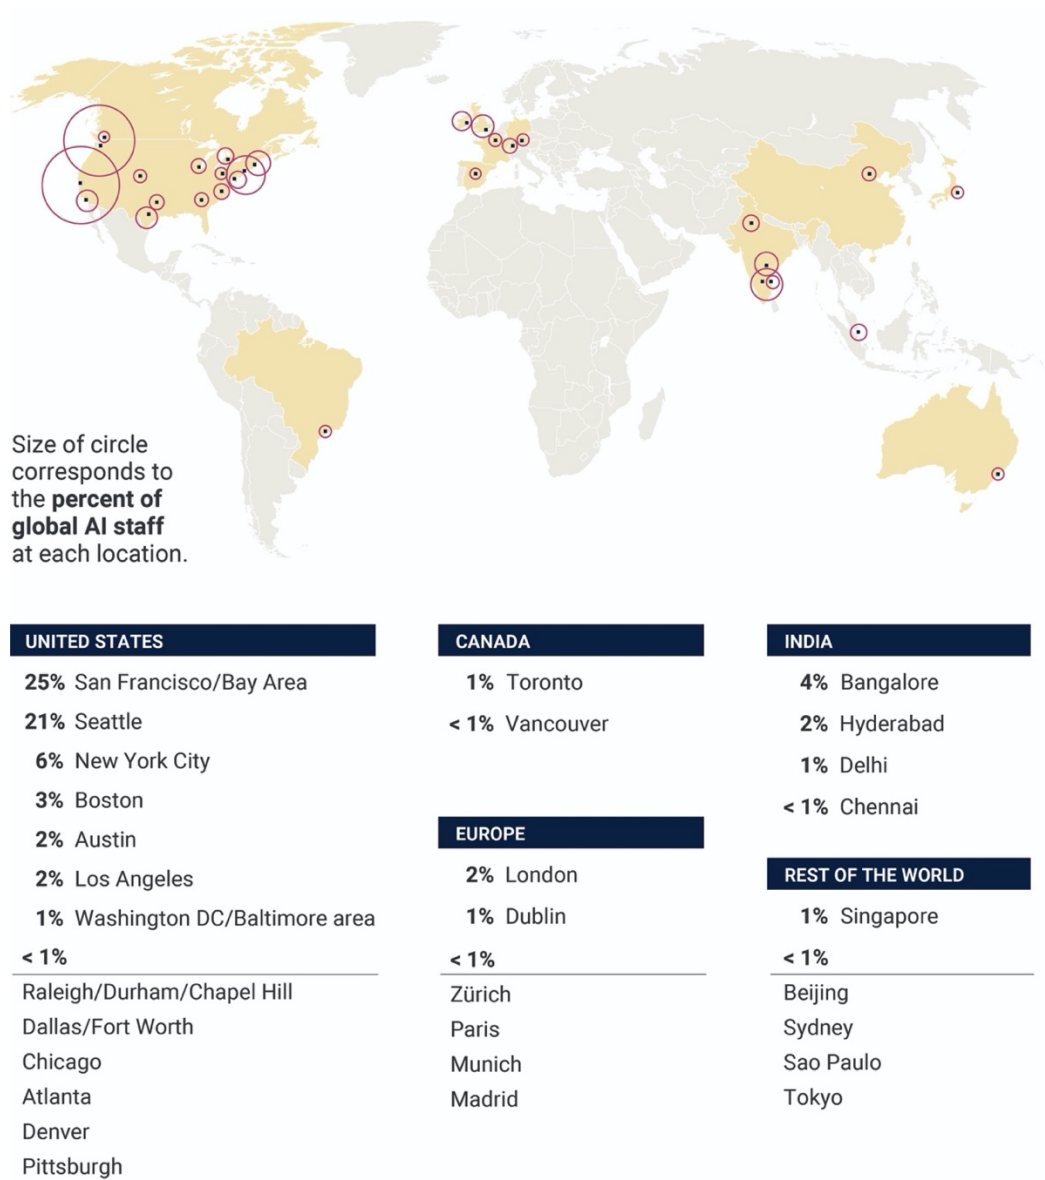
* One possible explanation for this finding is that Facebook and Google have set up many relatively new AI-specific research entities, while IBM and Microsoft only have general R&D facilities, some of which have recently added AI components. Other explanations, such as market-driven location choices, are discussed in more detail in the section of this paper on “What Drives Location Decisions.”

The previous two figures tracked the number and location of labs, but most companies do not release public data on how large their labs are. Moreover, as noted above, Amazon does not structure its AI R&D in labs and Apple does not publicly release information on their R&D activities, and the remaining four companies may do AI R&D outside of their labs as well. It is therefore important to compare these findings to staff counts, which provide another indicator of AI activity.

It was clear from Figure 3 that 68 percent of our companies' AI staff are located in the United States. Digging deeper, Figure 6 maps out all 30 metropolitan areas in 15 countries where our six companies cumulatively employ 500 or more AI staff.* Nearly half (46 percent) of AI staff are located in just two metropolitan areas: San Francisco and Seattle. Four other U.S. cities—Austin, Boston, Los Angeles, and New York—are also in the top ten locations.³²

* Locations with 500 staff represent locations that host more than 0.5 percent of these companies combined global workforces.

Figure 6. Locations where companies cumulatively have more than 500 AI staff.



Source: LinkedIn. Companies whose staff are included are Amazon, Apple, Facebook, Google, IBM, and Microsoft. See Appendix B for how “AI staff” was measured.

Outside of the United States, our six companies employ most AI staff across Europe, India, and the Asia-Pacific region. Several locations host AI labs but appear to have fewer than 500 AI staff, including several locations in China (in or near Shanghai, Shenzhen, and Beijing) and Israel. Because our main

source for AI staff is LinkedIn, these findings have to be interpreted with caution; some of the differences may be explained by different rates of LinkedIn usage among multinational employees in these different countries.*

What Types of R&D Do AI Labs Conduct?

This section looks at the types of R&D that companies conduct at their AI labs. R&D activities can be categorized as basic research, applied research, or experimental development.³³ In our context, we define “basic research” as the pursuit of foundational breakthroughs in AI; “applied research,” as the application of AI research to known problems; and “experimental development,” as the creation of AI prototypes and products.³⁴ In order to group the R&D activities of AI labs into these three categories, we looked at news coverage, press releases, and other corporate statements (see


* LinkedIn data likely over-represents AI staff from some nations relative to others, as the platform is far more common in, say, the United States than it is in China. We expect cross-country differences in the rates of LinkedIn usage to be significantly smaller among employees at U.S. multinationals than among populations at large, but the possibility of error remains real. For one study, which finds LinkedIn penetration rates ranging from 44 percent of the population in the United States to 1 percent in Nigeria, see Bloom, “LinkedIn country-by-country audience visualized.”

Appendix D for more detail and examples). Many labs conduct multiple types of R&D.*

Figure 7 shows the proportion of AI labs conducting each type of R&D. About 80 percent of the AI labs perform some form of research. Most AI labs (41 of 62) conduct some basic research, although AI labs often conduct more than one type of AI R&D. More than half of labs perform at least two types of AI R&D, corroborating a separate survey of multinational R&D labs in China (not specific to AI) that found “it is not unusual to observe all three types of R&D activities within one tech MNC R&D lab.”³⁵

* For example, Google Research India’s activities in Bangalore count as a mix of all three R&D types based on the following statement: “This team will focus on two pillars: First, advancing fundamental computer science and AI research ... [and] second, applying this research to tackle big problems ... while also using it to make apps and services”; quoted in Jay Yagnik, “Google Research India: an AI lab in Bangalore,” *Google: The Keyword* (blog), September 19, 2019, <https://blog.google/around-the-globe/google-asia/google-research-india-ai-lab-bangalore/>.


Figure 7. Types of R&D done at AI labs by R&D category.



Source: CSET AI labs data (see Appendix A).

Figure 8 shows the fraction of AI labs performing each type of R&D by country or region. Labs in North America and Europe are predominantly research focused, whereas labs located in other regions conduct less basic research and more experimental development. Google's AI China Center was, at least at the time of its launch, a notable exception; a press release announcing the center said it would be "focused on basic AI research."³⁶ Google's AI lab in Ghana also conducts both basic and applied research.³⁷


Figure 8. Types of R&D at AI labs by country or region.



Source: CSET AI labs data (see Appendix A). A single AI lab can be involved in multiple types of R&D. Australia and Taiwan, which each host only one lab, are excluded from the figure.

Figure 9 shows the fraction of each company’s AI labs conducting each type of R&D. There are significant differences between the companies in terms of the types of R&D their labs perform. Facebook AI Research (FAIR) labs are almost exclusively engaged in basic research, Google’s labs conduct a mix of basic and applied research, Microsoft’s labs are more likely to focus on experimental development, and IBM’s labs conduct a mix of all three R&D types.³⁸ This illustrates that tech companies structure their R&D activities differently, and may also use the term “lab” differently; as noted in Section 2, in this paper we adopted each company’s own “lab” designations. We discuss how we coded types of R&D in greater detail in Appendix D.

Figure 9. Types of AI R&D at labs by company.



Source: CSET AI labs data (see Appendix A). AI labs involved in multiple types of R&D are counted multiple times.

Differences in companies' AI R&D activities are also visible with respect to where they are established and how they performed. In Seattle, where AI labs primarily focus on basic and applied research, Facebook's local lab director retains a computer science professorship at the University of Washington and involves strong university ties.³⁹ By contrast, other companies' labs in Israel tend to focus on applied research and experimental development. For example, Microsoft AI labs there, such as the Israel Development Center and Advanced Technology Lab Israel, work closely with startup founders and have a stronger engineering focus.⁴⁰

What Drives Location Decisions?

This section considers why companies set up labs and offices in particular locations. We find that these decisions are typically made on a variety of grounds, but companies' statements, along with research in economics and business, suggest three main reasons: access to talent, market access or product adaptation, and cost savings.*

Access to Talent

Companies nearly always cite the presence of a local talent pool as a dominant reason for their location choices.⁴¹ Speaking about a series of new

* A fourth reason suggested by the business literature for foreign investment by multinationals is "technology scouting," finding good ideas or companies to collaborate with or acquire. Beyond Microsoft's labs in Israel—where Zack Weisfeld, the then general manager of Microsoft Ventures Global Accelerators, spoke of Advanced Technology Lab Israel as a means of finding promising AI startups—we did not see this motivation reflected in company statements. See Julie Bort, "After Success in Israel, Microsoft Launches More Accelerator Programs," *Business Insider*, April 24, 2013, <https://www.businessinsider.com/after-success-in-israel-microsoft-launches-worldwide-startup-accelerator-program-2013-4>. Moreover, we did not assess what effects U.S. export control policies have on companies' decisions to locate R&D overseas. Export controls may affect companies' choice of location—one IBM executive said export controls are "relevant" though "not a tier-one factor" in determining the company's R&D location choices; cited in Andrew B. Kennedy, *The Conflicted Superpower: America's Collaboration with China and India in Global Innovation* (New York: Columbia University Press, 2018), 135. However, export controls did not come up in news reports about AI labs, and the complexities of export control regulations make it difficult to assess independently.

Facebook AI labs, CTO Mike Schroepfer put it simply: "We're basically going where the talent is."⁴² A Google spokesperson said, "the primary driver for our choice of AI research office locations is the availability of talent who are either already living in the area or would be attracted to move there."⁴³ The managing director of Microsoft Research Asia explains the company's decision to open a lab in China as resulting from their realization that "we cannot just have the research lab in the US because we know we won't hire enough, and also, we don't believe all the smart and the most passionate people doing research will all come to [the United States]."⁴⁴ These arguments are reinforced by companies' perception that there is a global shortage of people with the necessary AI skills, which makes locating in talent clusters especially attractive.⁴⁵

Labs and offices are often located near universities; in the AI labs database, 38 percent of regions with at least one AI lab are also home to one or more of the world's top 50 AI universities.⁴⁶ Companies' local offices often forge close ties with universities in an attempt to access talent (see also Box 1). Examples include Google's ties with the Technical University of Munich and Microsoft's relationship with the University of Cambridge.⁴⁷ In Israel, Google reportedly sources half of its engineers from Tel Aviv University.⁴⁸ Companies are also often interested in hiring professors and researchers employed by universities⁴⁹—in one prominent example, Uber hired about 50 staff from Carnegie Mellon University in 2015 shortly after setting up a Pittsburgh lab.⁵⁰ Universities also incubate startups that tech companies want to keep an eye on for potential partnerships or acquisitions.

Some locations are good convening points for talent from a broader region, either because they are up-and-coming hubs or because they have friendly immigration policies. Google AI Ghana and Microsoft Taiwan, for example, appear to have been set up to recruit not just from their respective countries but from their broader regions.⁵¹ Labs within the European Union, where internal labor mobility is unrestricted for citizens, also often serve as hubs for regional talent.⁵²

Box 2. The Role of U.S. Immigration Policy in Location Decisions

Media reports and economic studies suggest U.S. immigration restrictions are associated with U.S. companies' global expansion. For instance, one July 2020 article on an executive order halting temporary employment-based programs said that "even before the new visa restrictions, immigration obstacles had begun to hurt AI activity in the U.S.," citing as evidence that "Companies like Facebook, Microsoft, Google, Amazon, and Intel established AI centers in other countries in pursuit of local talent."⁵³ IBM says that 58 percent of AI researchers in its U.S. labs were born abroad.⁵⁴

In line with this argument, several economic studies have found that immigration restrictions lead to tight labor markets for software talent in the United States, and that tight labor markets in turn drive companies to locate more of their R&D abroad in places where skilled workers are more plentiful.⁵⁵ In recent years, several countries have reformed their immigration systems to make it easier to attract AI talent.⁵⁶ Some have also launched recruitment campaigns that explicitly target U.S.-based workers with immigration difficulties, as illustrated by a Canadian billboard in Silicon Valley that reads: "H-1B Problems? Pivot to Canada."⁵⁷ A Toronto-based immigration lawyer found that U.S. tech firms "have opted to open Canadian offices [because they] can attract and keep workers in Canada."⁵⁸ A Microsoft Research representative told us, "some of the people we are hiring today in China and India are the exact same people we would normally be hiring in Redmond, Boston, or NYC, but today they are not able to get visas to immigrate to the U.S."⁵⁹

On the other hand, as discussed in this section, companies have compelling reasons beyond concerns with U.S. immigration policy to establish themselves abroad. For example, Facebook's AI lead Yann LeCun noted that "immigration in Canada is more well organized than immigration into the U.S.," but that "this didn't factor into [Facebook's] decision" to establish itself in Montreal.⁶⁰

In sum, the evidence suggests that U.S. companies would likely establish AI R&D operations abroad at a lower rate if U.S. immigration policies were less restrictive. However, due to factors beyond talent that drive location decisions, companies would continue conducting at least some of their AI R&D abroad even if immigration reforms were adopted.

Market Access and Product Adaptation

Economists often contrast "knowledge-seeking" or "talent-seeking" investments abroad with "market-driven" investments.⁶¹ Foreign labs can help gain access to markets in one of two ways. First, some countries require

companies to establish local facilities or partnerships in exchange for market access. Even in countries without such requirements, however, companies often set up local offices in order to be close to their customers and to adapt products to local conditions.⁶²

Market access does not appear to be a dominant motivation for the companies establishing new standalone AI labs in our dataset, perhaps because we exclude MNC offices without an R&D component (e.g., those exclusively staffed by product and sales teams). But market access may matter to AI R&D location indirectly. In China, for example, one study notes that “most of the handful of R&D centers that conduct basic research did not begin to do so until they had already established strong product development functions.”⁶³ Other studies find similar local “value chain climbing” in countries with sufficient R&D talent pools, including India and China.⁶⁴

Cost Savings

Companies also conduct R&D abroad in order to save on costs. It is commonly accepted that AI talent is expensive; exact numbers vary across studies, but most find average salaries well above \$100,000 in the United States.⁶⁵ Similarly detailed AI salary statistics are not available for other countries, but a 2013 survey by the National Science Foundation found that U.S. information and communication technologies companies on average had to spend nearly twice as much per employee for R&D conducted domestically compared to R&D performed abroad.⁶⁶ This indicates cost savings on salaries could be significant.⁶⁷

Host governments also sometimes attract multinational companies through non-salary incentives, for example by providing low-rent or rent-free office space or subsidized electricity. A desire to avoid paying U.S. taxes, or to take advantage of tax incentives provided by the host country, can also be motivators for companies. Host governments provide such incentives because they believe their populations, businesses, and economies will benefit. A host country’s eagerness for foreign technology can also impose costs on multinationals, however. These costs include the risk of physical or intellectual property theft and the expense of guarding against such risks.

Companies themselves are typically reluctant to admit cost plays a role in their decisions to establish offices abroad, but experts recognize it as an important factor.⁶⁸ For example, evaluating Microsoft’s establishment in

Taiwan, which included an AI R&D component, a Taipei-based industry analyst pointed to the greater return on investment the company would have in Taiwan due to lower salaries.⁶⁹ Analysts also argue, however, that lower costs, while helpful, are not a sufficient reason for a company to open a new location. Without a supply of sufficiently high-quality local talent, companies will be reluctant to launch initiatives involving significant technical work even if costs are substantially lower.⁷⁰

Conclusion

Based on our analysis of new data on U.S. multinational companies' AI R&D labs, we found that Facebook, Google, IBM, and Microsoft have 62 labs that are spread all over the world; 20 are in the United States, and the remaining 42 are located across 16 countries. In contrast to AI labs, however, most company AI staff remain concentrated in the United States. Much AI R&D in North America and Europe appears focused on research, whereas labs based elsewhere tend to emphasize development. Companies decide to conduct AI R&D abroad for multiple reasons, but access to global talent is typically cited as the driving factor. Our findings regarding AI labs are based on company classifications of "AI" and "labs"—we did not independently assess the work that is being done at these locations. Data on AI staff comes from information provided by individuals through LinkedIn.

Our findings can inform policy on several fronts. First, they reinforce past research that suggests U.S. immigration reforms would encourage U.S. companies to conduct more R&D at home (see Box 2), although it appears unlikely that increasing the domestic talent supply would entirely halt or reverse companies' internationalization of R&D. Second, for other tools that could be used to encourage more domestic R&D, our data can help policymakers assess how different policy approaches would affect leading U.S. AI companies. Such impact assessments are an important part of balancing the complex benefits and risks of R&D internationalization. Third, calls for multilateral R&D initiatives and an "alliance innovation base" recommend building on pre-existing R&D links.⁷¹ Our data provides a window into places where U.S. companies already have a significant AI R&D presence, which could inform future diplomatic engagements.

This study also provides lessons and resources relevant to analysts. For example, it is important to look at multiple metrics of AI R&D abroad in studying the extent of U.S. multinationals' R&D internationalization, as illustrated by the fact that domestic-to-foreign ratios for AI labs and AI staff are almost mirror images. In Appendix C, we discuss several other AI R&D metrics that future studies should consider. In order to facilitate future work on


these questions, we are releasing our AI labs dataset along with this paper.* We encourage other researchers to build on this dataset by adding further information on the AI labs or by adding additional labs or companies. Other U.S. companies with AI R&D labs include Intel, NVIDIA, and Uber,⁷² and many non-U.S. companies conduct significant global AI R&D, as well.⁷³

* The CSET AI labs dataset is available at <https://github.com/georgetown-cset/us-multinational-ai-labs>.

Acknowledgments

Thanks to Tessa Baker, Ben Chang, Melissa Flagg, Rebecca Kagan, Igor Mikolic-Torreira, Dewey Murdick, Ilya Rahkovsky, Kathleen Walsh, and representatives from Facebook, Google, and Microsoft for their input and feedback; to Rebecca Gelles and Farhana Hossain for data and visualization support; and to Dale Brauner, Melissa Deng, Shelton Fitch, and Lynne Weil for editorial support.

Roxanne Heston co-authored this paper while she was a Research Analyst at CSET. She is currently detailed as an artificial intelligence policy adviser in the Office of the Science and Technology Adviser at the U.S. Department of State. She completed her contributions to this paper prior to joining the State Department. The views expressed herein are the authors' and do not necessarily reflect those of the U.S. government. Inquiries about this report should be directed to Remco Zwetsloot.



© 2020 by the Center for Security and Emerging Technology. This work is licensed under a Creative Commons Attribution-Non Commercial 4.0 International License.

To view a copy of this license, visit
<https://creativecommons.org/licenses/by-nc/4.0/>.

Appendix: Data and Methods

A. Locating AI Labs

To locate AI labs, we looked primarily at company news releases, company websites, and other news coverage. Labs could involve either a standalone AI R&D group or an AI R&D-focused team within a larger R&D organization.* When companies listed R&D facilities that were not explicitly named as “AI” centers, we looked at their websites to determine whether they described themselves as having a systematic AI R&D program. We relied on companies’ own choices for what they called “AI” and what they labeled a “lab.”

This approach has several shortcomings. First, some information was not available for all labs. For example, we were unable to determine the dates of establishment for a majority of the AI labs in our dataset. More broadly, media and companies themselves often discuss projects outside of the United States in more detail because they are considered more notable, though some labs abroad received little coverage in English-language sources. More detailed surveys of the companies may be needed to fill these informational gaps.

Second, what counts as a “lab” differs across companies, and two of the companies we studied, Amazon and Apple, either do not use or do not publicize “labs” (see Section 2).⁷⁴ Designations are often a form of branding.

* The following companies listed R&D locations on their websites. Only Facebook’s AI labs were consistently listed distinct from their general R&D facilities:

- Facebook: <https://web.archive.org/web/20190611075902/https://research.fb.com/category/facebook-ai-research/> (The live version of this link now redirects to ai.facebook.com, but a member of Facebook’s AI Communications team corroborated this list in an email exchange.)
- Google: <https://research.google/careers/>
- IBM: <https://www.research.ibm.com/labs/>
- Microsoft: <https://www.microsoft.com/en-us/research/>

A Google representative, for example, explained that the company used to designate some of its locations as “AI centers,” but that groups doing the same work today are simply referred to as Google Research locations.⁷⁵ Similarly, Microsoft Research emphasizes that “AI research is infused across many parts of Microsoft,” and does not happen only within groups designated as “labs.”⁷⁶ This shortcoming can be addressed by looking at alternative metrics of R&D. This paper also looked at AI staff; Appendix C lists several additional AI R&D-related metrics that future research could collect data on.

We hope future research will shed further light on these issues. To lay the groundwork for follow-on research, we are releasing our AI labs dataset along with this paper. In addition to the descriptive AI labs data that we reported on in the paper, the dataset also includes links to the sources that we used to make coding decisions, for example on the type(s) of R&D that each lab is engaged in. We did not conduct site visits at labs or interview lab leadership and staff, host country governments, U.S. government experts, or other stakeholders. Future research should also look beyond quantitative metrics to include such visits and interviews, if possible.

B. Locating AI Staff

To count the number of “AI staff” at our six companies across various locations, we used employees’ skills listed on their LinkedIn profiles. We accessed LinkedIn data through the LinkedIn Talent Insights platform. To define “AI” skills, we used a list of “core AI” skills supplied to CSET by Burning Glass Analytics, a labor market analytics company. In total, there are 46 “AI skills” that were included in both Burning Glass’s AI taxonomy and found on LinkedIn profiles, ranging from specific machine learning libraries (e.g. PyTorch, TensorFlow) to general skill categories (e.g. Data Science, Machine Learning).*

* The following 46 skills both appear on LinkedIn and were identified as “core AI skills” by Burning Glass: Artificial Intelligence (AI), IBM Watson, iThink, Cluster Analysis, Data Science, k-means clustering, Kernel, Pattern Recognition, Predictive Analytics, Principal Component Analysis, Computer Vision, Decision Trees, Deep Learning, Deeplearning4j,

Any binary measure of whether someone has “AI skills” will inevitably be imperfect. Narrow operationalizations will mean excluding people who do AI R&D work (“false negatives”), whereas a broad operationalization will apply to some people without real AI skills (“false positives”). On balance, our skill list leans toward a broad operationalization; to minimize false negatives, we include skills like “Chatbot” that may capture some workers who are not conducting AI R&D. Moreover, using LinkedIn profiles means that we are relying on people to accurately report their skills. Finally, as discussed in Sections 2 and 3, LinkedIn’s differential coverage may also result in undercounting in countries outside of North America or Europe. These issues should be kept in mind when interpreting the results.

Researchers also have to determine how to associate employees with companies. Multinational corporations frequently have many wholly owned subsidiaries or affiliates, such as Amazon Web Services for Amazon. In measuring staff counts in particular locations, researchers have to choose whether to include only employees at the main company or also at their subsidiaries and affiliates. In Figures 2, 3, and 6 we reported results only for the main company. As a robustness check we also reran our analysis including affiliates, which did not significantly change the results.⁷⁷

Pending data availability, future research could use staff-based measures to dig deeper on several questions. For example, if demographic data on staff were available, future work could analyze what percentage of staff at U.S. and foreign locations are U.S. citizens, and what the most common non-U.S. nationalities are among company employees across different locations. Similarly, data on staff job titles and educational backgrounds could, as discussed in Appendix C below, be used to assess the type(s) of AI work being done in different locations.

Google Cloud Platform (GCP), Gradient Boosting, H2O.ai, Keras, Machine Learning, Apache Mahout, Microsoft Cognitive Toolkit (CNTK), MXNet, Neural Networks, Object Tracking, OpenCV, Random Forest, Recommender Systems, Support Vector Machine (SVM), Vowpal Wabbit, XGBoost, Caffe, TensorFlow, PyTorch, ANTLR, Chatbots, Computational Linguistics, Latent Dirichlet Allocation, Lexical Semantics, Machine Translation, MoSes, Natural Language Processing (NLP), K-Nearest Neighbors (KNN), Sentiment Analysis, Speech Recognition, Text-to-Speech, Tokenization.

C. Comparing AI Labs, AI Staff, and Other R&D Measures

This paper included data on both AI labs and AI staff because looking at only a single metric is likely to provide a partial and misleading picture.⁷⁸ Looking only at labs, for example, could make it seem like the bulk of U.S. MNC AI R&D is happening abroad, while looking at staff could give the opposite impression. This does not mean the two metrics are contradictory; both have strengths and weaknesses, and both tell an important story.

Where a multinational launches AI labs, and the company statements and media coverage that accompany lab launches, can provide a window into that company's AI R&D strategy. These statements often convey the company's expansion plans and the type of R&D they hope to be doing, and explain why the company decided to open the lab where it did. But company statements must be taken with a grain of salt, so there are limits to what can be gleaned from this type of information. Augmenting labs data with other metrics is also important because not all companies use "labs" to structure their R&D investments.

Measures of staff counts provide a more objective indicator of how much a company has invested in a specific location. Companies might be inclined to say that all labs are very valuable to them, but a location with 200 staff is likely to be more valuable than one with 20 staff. A weakness of current staff measures, however, is that they only provide a snapshot: a location with 20 staff could have 200 staff one or two years later. Depending on how much data is available on staff—do researchers only have data on raw numbers or also on job titles, professional background, and so forth?—staff data may or may not allow a detailed assessment of a company's focus. In this paper we only reported raw staff numbers, but future research could go deeper.

Future research on the global distribution of AI R&D by multinational companies can also look at other metrics. Additional metrics could include:⁷⁹

- Job advertisements. In addition to overall staff counts, job advertisements can provide an indicator of a company's growth in a specific location. It might be possible to use advertisement content, such as educational requirements or job titles, to distinguish AI jobs from non-AI jobs and R&D jobs from non-R&D jobs.

- Spending. Company figures on AI R&D spending across different locations would provide a valuable look into their R&D activity. However, companies may be unable or unwilling to share data that looks at a specific technology area (AI) and disaggregates by location. Existing surveys that ask companies for R&D spending, such as the National Center for Science and Engineering Statistics' Business R&D and Innovation Survey (BRDIS), unfortunately do not contain sufficiently detailed data for AI-specific analysis.⁸⁰
- Acquisitions. Large multinationals frequently acquire startups and other companies in specific locales, sometimes to set up R&D labs (see Box 1). One possible issue with this metric from an R&D perspective is that it is not always clear whether an acquisition was focused on R&D activities.
- Partnerships. Companies often partner with other organizations, including universities, in setting up research activities. For example, Microsoft, as part of its \$1.65 billion "AI for Good" initiative, provides extensive support to the University of Waterloo's AI Institute.⁸¹ Such partnerships were excluded from our dataset but are part of companies' local AI R&D footprints. Partnership data is currently not tracked in a systematic fashion.
- Publications. Looking at papers that staff from particular company labs publish could provide another indicator of AI R&D activity. Possible shortcomings include limited data availability on specific locations (authors may not report their geographic location in their papers), poor coverage of more applied work (which is less likely to get published), and systematic differences in company policies that would make cross-company comparisons problematic (some companies allow or encourage their staff to publish whereas others do not).
- Patents. Company patents could also function as an AI R&D indicator, and inventor addresses can help link the employees to specific company locations.⁸² For example, in a press release, IBM said that "IBM inventors who reside outside the U.S. contributed to more than 36 percent of the company's 2015 patents."⁸³ As with publications, however, patents only cover a subset of R&D work, and cross-

company comparisons could be problematic due to policy differences.⁸⁴

Future research should also look at the nature and impact of foreign R&D activities more broadly. For example, to what extent are U.S. multinationals able to control access to the R&D done in their labs abroad? How are property rights and ownership divided between multinationals and local partners across different projects? How do companies weigh national security concerns when establishing locations abroad?

D. Classifying AI R&D Activities

In the section on “What Types of R&D Do AI Labs Conduct,” we reported results that categorized R&D into three categories: (1) Basic research; (2) Applied research; and (3) Experimental development. Categories were not exclusive, i.e., one lab can conduct multiple types of R&D. Below we provide some examples, drawn from company websites and news reports, to illustrate the type of information on which we based our coding choices. The CSET AI R&D labs dataset released alongside this paper contains hyperlinked sources for each choice.

- Basic research: “Members of FAIR carry out fundamental research in the field of AI. Some of their breakthroughs are applied to Facebook’s platforms (Facebook, Instagram, and WhatsApp) by an applied machine learning (AML) team and other engineers, but the majority of their research is purely academic.”⁸⁵
- Applied research: “Established in 1982, IBM Research – Tokyo focuses on innovations that use cognitive computing to solve social and industry problems through expertise in cognitive device technologies, text analytics and mathematical science technologies. With facilities in Tokyo and Shin-Kawasaki, IBM Research – Tokyo has been playing a key role in exploring research with clients in Japan by applying research assets and expertise to meet their needs for innovations and to overcome their business challenges through joint collaboration.”⁸⁶
- Experimental development: “The Microsoft AI R&D Center in Taiwan will not pursue the basic science of artificial intelligence, which means

teaching machines to interpret and anticipate a user's movements. Instead, the center is designed to be more engineering-centric."⁸⁷

- Mix of all three R&D categories: "This [Google Research India] team will focus on two pillars: First, advancing fundamental computer science and AI research by building a strong team and partnering with the research community across the country. Second, applying this research to tackle big problems in fields like healthcare, agriculture, and education while also using it to make apps and services used by billions of people more helpful."⁸⁸

Endnotes

¹ Jacques Bughin, Eric Hazan, Sree Ramaswamy, Michael Chui, Tera Allas, Peter Dahlström, Nicholas Henke, and Monica Trench, “Artificial Intelligence: The Next Digital Frontier?” (McKinsey Global Institute, June 2017), 10, <https://www.mckinsey.com/~media/McKinsey/Industries/Advanced%20Electronics/Our%20Insights/How%20artificial%20intelligence%20can%20deliver%20real%20value%20to%20companies/MGI-Artificial-Intelligence-Discussion-paper.ashx>. This is the most recent year for which we were able to find estimates of private sector R&D spending on AI.

² “The Race For AI: Here Are The Tech Giants Rushing To Snap Up Artificial Intelligence Startups,” CB Insights, September 17, 2019, <https://interactives.cbinsights.com/artificial-intelligence-acquisitions-by-famga/>.

³ President’s Council of Advisors on Science and Technology, *Recommendations for Strengthening American Leadership in Industries of the Future* (Washington, DC: Office of Science and Technology Policy, June 2020), 5, https://science.osti.gov/-/media/_pdf/about/pcast/202006/PCAST_June_2020_Report.pdf?la=en&hash=019A4F17C79FDEE5005C51D3D6CAC81FB31E3ABC.

⁴ Of that \$5 billion, \$4 billion is for AI-related research, development, testing, and evaluation (RDT&E) at the Department of Defense, and \$973.5 million is for civilian agencies; see Chris Cornillie, “Finding Artificial Intelligence Money in the Fiscal 2020 Budget,” Bloomberg Government, March 28, 2019, <https://about.bgov.com/news/finding-artificial-intelligence-money-fiscal-2020-budget/>; National Science and Technology Council, *Supplement to the President’s FY2020 Budget* (Washington, DC: Executive Office of the President of the United States, 2019), 7, <https://www.whitehouse.gov/wp-content/uploads/2019/09/FY2020-NITRD-AI-RD-Budget-September-2019.pdf>.

⁵ Rani Molla, “Amazon spent nearly \$23 billion on R&D last year — more than any other U.S. company,” *Vox*, April 9, 2018, <https://www.vox.com/2018/4/9/17204004/amazon-research-development-rd>.

⁶ Omri Wallach, “How Tech Giants Make Their Billions,” *Visual Capitalist*, July 6, 2020, <https://www.visualcapitalist.com/how-big-tech-makes-their-billions-2020/>; Vaishali

Ladha, "IBM: Revenue by Region (2020)," Business Quant, <https://businessquant.com/ibm-revenue-by-region-2>.

⁷ Gary Clyde Hufbauer, Theodore H. Moran, Lindsay Oldenski, and Martin Vieiro, *Outward Foreign Direct Investment and US Exports, Jobs, and R&D: Implications for US Policy*, Peterson Institute for International Economics, August 2013, <https://www.piie.com/bookstore/outward-foreign-direct-investment-and-us-exports-jobs-and-rd-implications-us-policy>.

⁸ Department of Defense, *Summary of the 2018 National Defense Strategy* (Department of Defense, 2018), 3, <https://www.hsdl.org/?view&did=807329>.

⁹ Darren E. Tromblay and Robert G. Spelbrink, *Securing U.S. Innovation: The Challenge of Preserving a Competitive Advantage in the Creation of Knowledge* (New York: Rowman & Littlefield, 2016), 68; William C. Hannas and Huey-meei Chang, "China Access to Foreign AI Technology" (Center for Security and Emerging Technology, September 2019), 13, <https://cset.georgetown.edu/research/chinas-access-to-foreign-ai-technology/>; Lee Branstetter, Britta Glennon, and J. Bradford Jensen, "Knowledge Transfer Abroad: The Role of U.S. Inventors within Global R&D Networks" (NBER Working Paper No. 24453, March 2018), <https://www.nber.org/papers/w24453>.

¹⁰ Matt Sheehan, "Who Benefits From American AI Research in China?", MacroPolo, October 21, 2019, <https://macropolo.org/china-ai-research-resnet/>. On a webpage about its presence in China, Microsoft says that it "boasts a robust partner ecosystem with 17,000 partners. For every RMB that Microsoft earns in China, Microsoft partners earn 16"; see "About Microsoft's Presence in China," <https://news.microsoft.com/about-microsofts-presence-in-china/>.

¹¹ Ryan Browne, "Top US general says Google 'is indirectly benefiting the Chinese military,'" *CNN Politics*, March 14, 2019, <https://www.cnn.com/2019/03/14/politics/dunford-china-google/>.

¹² A Microsoft Research representative told us that Microsoft believes "China would be where it is now regardless of what we did in the early 2000s in China" (email correspondence, October 6, 2020). Google's CEO has contested claims the company benefits the Chinese military and says the company has "cleared up any misunderstandings"; see Jackie Wattles, "Google CEO says he wants to reach the next billion users, but has no plans to relaunch in China," *CNN*, June 17, 2019, <https://edition.cnn.com/2019/06/16/tech/google-china-ceo-sundar-pichai/index.html>.

¹³ See, e.g., White House, "National Security Strategy," December 2017, 20, <https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905-2.pdf> (calls for U.S. government agencies to "improve their understanding of worldwide S&T trends and how they are likely to influence--or undermine--American strategies and programs"); Sheehan, "Who Benefits From American AI Research in China?"; Tromblay and Spelbrink, Securing U.S. Innovation.

¹⁴ On R&D spending, see Molla, "Amazon spent nearly \$23 billion on R&D last year — more than any other U.S. company"; on patent counts, see TR Staff, "Who Is Winning the AI Race?", *MIT Technology Review*, June 27, 2017, <https://www.technologyreview.com/s/608112/who-is-winning-the-ai-race/>; on startup acquisitions, see CB Insights, "The Race For AI: Here Are The Tech Giants Rushing To Snap Up Artificial Intelligence Startups"; on industry publications, see Gleb Chuvpilo, "Who's Ahead in AI Research? Insights from NIPS, Most Prestigious AI Conference," *Medium*, August 7, 2018, <https://medium.com/@chuvpilo/whos-ahead-in-ai-research-insights-from-nips-most-prestigious-ai-conference-df2c361236f6>; on faculty hiring, see Michael Gofman and Zhao Jin, "Artificial Intelligence, Human Capital, and Innovation," SSRN, January 2020, 36, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3449440. All companies are prominent members of the Partnership on AI; see "Meet the Partners," Partnership on AI, <https://www.partnershiponai.org/partners/>.

¹⁵ Steven Levy, "Inside Amazon's Artificial Intelligence Flywheel," *WIRED*, February 1, 2018, <https://www.wired.com/story/amazon-artificial-intelligence-flywheel/>; Amazon uses the term "Tech Hub" to describe some of the locations where the company conducts research and software development; see, e.g., "Amazon Expands Toronto Tech Hub and Announces Plans to Create 600 New Tech Jobs," *dayone: The Amazon Blog*, December 18, 2018, <https://press.aboutamazon.com/news-releases/news-release-details/amazon-expands-toronto-tech-hub-and-announces-plans-create-600>.

¹⁶ On Apple's relative secretiveness, see, e.g., Tristan Greene, "Apple's culture of secrets is why it won't win the AI race," *The Next Web*, September 6, 2017, <https://thenextweb.com/apple/2017/09/06/apples-culture-of-secrets-is-why-it-wont-win-the-ai-race/>. The public information on its AI R&D activities that exists suggests the company does have some AI R&D labs. For example, Apple's offices in Cambridge (United Kingdom) "are thought to be hubs for Siri and A.I. research," according to news reports; Jonny Evans, "Apple is pumping money into creating an international chain of R&D centers," *ComputerWorld*, March 17, 2017, <https://www.computerworld.com/article/3181020/apple-is-pumping-money-into-creating-an-international-chain-of-randd-centers.html>.

¹⁷ Google's lab predominance may be in part attributable to how they delineate AI labs. Google often has different names for different AI projects in the same market area. For instance, Google has three distinct AI labs in Montreal: Google Research, Google Brain, and DeepMind. Other companies do not as clearly separate their AI projects, which may drive up Google's lab count without necessarily indicating an increase in the number of people they have doing AI research in a location. However, these different AI labs do often focus on different types of R&D and may indicate a diversity of research that other companies do not have in any one location. See Appendix A for a discussion of how we defined AI labs.

¹⁸ Data limitations mean we do not have the number of employees by company, or company-specific location data.

¹⁹ LinkedIn data likely over-represents AI staff from some nations relative to others, as the platform is far more common in, e.g., the United States (44% penetration) than it is in China (3% penetration). However, we expect the disparity in representation by country is likely to be significantly lower for the subset of the population that works at U.S. technological MNCs. Jeremy Bloom, "LinkedIn country-by-country audience visualized," Thinknum, March 29, 2018, <https://media.thinknum.com/articles/tracking-linkedins-growth-from-american-icon-to-global-player/>.

²⁰ We compared the fraction each company's AI staff comprised with and without their affiliates, but the results stayed similar to those above. This graph shows the companies alone, affiliates excluded.

²¹ Romain Dillet, "Facebook Opens New AI Research Center In Paris," *TechCrunch*, June 2, 2015, <https://techcrunch.com/2015/06/02/facebook-opens-new-ai-research-center-in-paris/>; Sam Shead, "Facebook Plans To Double Size Of AI Research Unit By 2020," *Forbes*, October 1, 2018, <https://www.forbes.com/sites/samshead/2018/10/01/facebook-plans-to-double-size-of-ai-research-unit-by-2020/#32f8cf1c4c4f>.

²² Market areas are clusters of towns and cities that are located near each other and tightly economically integrated. For example, San Francisco and its surrounding towns (e.g., Menlo Park, Palo Alto) are counted as falling within the same market area. In practice, these market areas overlap with the OECD's metropolitan area classifications (<https://stats.oecd.org/Index.aspx?DataSetCode=CITIES>). For non-OECD countries, we applied the OECD's framework to that country's cities and regions; cities in Israel's "Silicon Wadi" region — which include Tel Aviv, Jerusalem, and Haifa — are counted as falling within the same market area because they are tightly economically integrated. For more discussion, see Organization for Economic Cooperation and Development, *Policy*

Roundtable: Market Definition, 2012,
<https://www.oecd.org/daf/competition/Marketdefinition2012.pdf>.

²³ For AI startup ecosystems, see “Global Startup Ecosystem Report 2019,” Startup Genome, May 9, 2019, 67-68, <https://startupgenome.com/reports/global-startup-ecosystem-report-2019>; for top AI programs, see Emery Berger, “CSRankings: Computer Science Rankings,” CSRankings, Accessed August 7, 2020, <https://csrankings.org/#/index?ai&vision&mlmining&nlp&ir&world>.

²⁴ For an overview of the economics literature on tech clusters, see William R. Kerr and Frédéric Robert-Nicoud, “Tech Clusters,” The National Bureau of Economic Research, no. 27421 (June 2020), <https://www.nber.org/papers/w27421>.

²⁵ Synced, “Google Boosting its AI Research in Tokyo,” *Medium*, April 27, 2018, <https://medium.com/syncedreview/google-boosting-its-ai-research-in-tokyo-608f8ba11c9>.

²⁶ Cade Metz, “Facebook Adds A.I. Labs in Seattle and Pittsburgh, Pressuring Local Universities,” *The New York Times*, May 4, 2018, <https://www.nytimes.com/2018/05/04/technology/facebook-artificial-intelligence-researchers.html>.

²⁷ “MIT-IBM Watson AI Lab,” <https://mitibmwatsonailab.mit.edu/>; Darrell Etherington, “Microsoft creates an AI research lab to challenge Google and DeepMind,” *TechCrunch*, July 12, 2017, <https://techcrunch.com/2017/07/12/microsoft-creates-an-ai-research-lab-to-challenge-google-and-deepmind/>.

²⁸ Microsoft Asia News Center, “Microsoft establishes Microsoft Research Asia-Shanghai and a new AI Innovation Center,” *Microsoft Blog*, September 17, 2018, <https://news.microsoft.com/apac/2018/09/17/microsoft-establishes-microsoft-research-asia-shanghai-and-a-new-ai-innovation-center/>.

²⁹ We know that five of the six companies made 59 (9.2 percent) AI acquisitions. We do not have data for IBM but, given how the data is presented, believe IBM made between one and six such purchases, for a total of 9.4-10.2 percent. CB Insights, “The Race For AI: Here Are The Tech Giants Rushing To Snap Up Artificial Intelligence Startups.”

³⁰ Microsoft News Center, “Microsoft announces significant expansion of Montreal research lab, new director,” *Microsoft News Centre Canada*, January 24, 2018,

<https://news.microsoft.com/en-ca/2018/01/24/microsoft-announces-significant-expansion-montreal-research-lab-new-director/>.

³¹ For example, when Google acquired Jelpac, an AI-based image recognition company, in 2014, their staff, intellectual property, and functions were rolled into Google's existing (now-defunct) Picasa product.

³² For a deeper look at the geographic distribution of AI talent in the United States beyond the six big tech companies, see Melissa Flagg and Justin Olander, "AI Hubs in the United States" (Center for Security and Emerging Technology, May 2020), <https://cset.georgetown.edu/research/ai-hubs-in-the-united-states/>.

³³ This classification is used, for example, by the U.S. National Science Board, which adopted the definitions from the Organization for Economic Cooperation and Development. For complete definitions, see the OECD glossary entries for Basic Research (<https://stats.oecd.org/glossary/detail.asp?ID=192>), Applied Research (<https://stats.oecd.org/glossary/detail.asp?ID=12>), and Experimental Development (<https://stats.oecd.org/glossary/detail.asp?ID=908>).

³⁴ The boundaries between these types of R&D are undoubtedly blurry; see, for example, Donald E. Stokes, *Pasteur's Quadrant* (Washington, DC: The Brookings Institution, 1997).

³⁵ Xiaohong Quan, "Multinational Research and Development Labs in China: Local and Global Innovation," Ph.D. dissertation, University of California, Berkeley, 2005, ProQuest (3211492).

³⁶ Fei-Fei Li, "Opening the Google AI China Center," *Google: The Keyword* (blog), December 13, 2017, <https://www.blog.google/topics/google-asia/google-ai-china-center/>.

³⁷ Google's two AI labs in Ghana are Google AI Ghana and Google Brain Ghana, both located in the capital city Accra. Google AI Lead Jeff Dean said that they chose Ghana in 2019 for its "robust network of academic institutions as well as infrastructure." Victor Asemota, "'Ghana is the future of Africa': Why Google built an AI lab in Accra," *CNN World*, July 15, 2018, <https://edition.cnn.com/2018/07/14/africa/google-ghana-ai/index.html>.

³⁸ The one group that Facebook calls a "lab" and that does AI experimental development is Data.AI, a 10-person team improving the company's internal interface at Facebook's Tel

Aviv R&D center. CTech, "Facebook Sets Up AI Team in Tel Aviv," CTech, April 23, 2019, <https://www.calcalistech.com/ctech/articles/0,7340,L-3760927,00.html>.

³⁹ Yann LeCun, "Facebook AI Research Expands With New Academic Collaborations," *Facebook Newsroom*, July 17, 2018, <https://about.fb.com/news/2018/07/facebook-ai-research-expands/>.

⁴⁰ "Advanced Technology Lab Israel," Microsoft Research, <https://www.microsoft.com/en-us/research/group/advanced-technology-lab-israel/>; Microsoft Israel R&D Center, <https://www.microsoftrd.co.il/>; Sam Shead, "Execs from Facebook, Google, and Microsoft explain why they use Israel for their R&D," *Business Insider*, October 6, 2016, <https://www.businessinsider.com/facebook-google-microsoft-israel-rd-2016-10?op=1>.

⁴¹ Bronwyn H. Hall, "The Internationalization of R&D" (University of California, Berkeley, August 2010), 10, https://eml.berkeley.edu/~bhhall/papers/BHH10_RND_international_August.pdf.

⁴² Cade Metz, "Facebook Adds A.I. Labs in Seattle and Pittsburgh, Pressuring Local Universities," *The New York Times*, May 4, 2018, <https://www.nytimes.com/2018/05/04/technology/facebook-artificial-intelligence-researchers.html>.

⁴³ A Google spokesperson said: "In general, the primary driver for our choice of AI research office locations is the availability of talent who are either already living in the area or would be attracted to move there. Aside from that it is the usual factors, such as strength of local academic institutions, business-friendly environment and good transport infrastructure. We think that with world-class universities and a thriving developer community, Europe is well positioned to play a leading role in AI research and application." See Sam Shead, "Google Brain Has Expanded To Amsterdam," *Forbes*, July 10, 2018, <https://www.forbes.com/sites/samshead/2018/07/10/google-brain-has-expanded-to-amsterdam/#5e7f426b26f4>.

⁴⁴ Microsoft Research Podcast, "Celebrating 20 Years of MSR in Asia with Dr. Hsiao-Wuen Hon," <https://www.microsoft.com/en-us/research/podcast/celebrating-20-years-of-msr-in-asia-with-dr-hsiao-wuen-hon/>.

⁴⁵ On perceptions of talent shortages in AI, see Remco Zwetsloot, Roxanne Heston and Zachary Arnold, "Strengthening the U.S. AI Workforce: A Policy and Research Agenda"

(Center for Security and Emerging Technology, September 2019),
<https://cset.georgetown.edu/research/strengthening-the-u-s-ai-workforce/>.

⁴⁶ Calculated using CSRankings top 50 from August 7, 2020.

⁴⁷ TU Munich's Applied AI initiative is designed to benefit Google and Cambridge's staff work on major Microsoft projects and its AI Residency Program; see, respectively, Arne Graevemeyer, "Google and TU Munich promote artificial intelligence [Google und TU München fördern Künstliche Intelligenz]," *Heise Online*, February 16, 2018, <https://www.heise.de/newsticker/meldung/Google-und-TU-Muenchen-foerdern-Kuenstliche-Intelligenz-3972112.html>; and Christopher Bishop, "Microsoft Research and Cambridge University strengthen their commitment to AI innovation and the field's future leaders," *Microsoft News Centre UK*, October 31, 2018, <https://news.microsoft.com/en-gb/2018/10/31/microsoft-research-and-cambridge-university-strengthen-their-commitment-to-ai-innovation-and-the-fields-future-leaders/>.

⁴⁸ Shead, "Execs from Facebook, Google, and Microsoft explain why they use Israel for their R&D."

⁴⁹ Facebook's Chief AI Scientist Yann LeCun says the company has set up AI research labs near universities in order to accommodate professors who do not want to move. "Basically you have to get the talent where it is," said LeCun. "Not everyone wants to live wherever we have our labs." Quoted in Alan Boyle, "FAIR competition? Facebook creates official AI labs in Seattle and Pittsburgh, vying for top talent," *GeekWire*, May 5, 2018, <https://www.geekwire.com/2018/fair-competition-facebook-raises-status-ai-research-labs-seattle-pittsburgh/>.

⁵⁰ Josh Lowensohn, "Uber gutted Carnegie Mellon's top robotics lab to build self-driving cars," *The Verge*, May 19, 2015, <https://www.theverge.com/transportation/2015/5/19/8622831/uber-self-driving-cars-carnegie-mellon-poached>.

⁵¹ Asemota, "'Ghana is the future of Africa'"; Elvina Yang, "Microsoft R&D Center in Taiwan Starts Recruiting for AI Research," *Meet Startup @TW*, March 27, 2018, <https://meet.bnext.com.tw/intl/articles/view/42604>.

⁵² McKinsey, for example, noted that Facebook's AI lab in Paris will "make it easier for the company to recruit top researchers in Europe." Bughin, Hazan, Ramaswamy, Chui, Allas, Dahlström, Henke, and Trench, "Artificial Intelligence: The Next Digital Frontier?"

⁵³ Kyle Wiggers, "White House advisory council's AI guidance conflicts with Trump's talent pool sabotage," *VentureBeat*, July 10, 2020, <https://venturebeat.com/2020/07/10/white-house-advisory-council-calls-on-u-s-to-increase-ai-funding-to-10-billion-by-2030/>.

⁵⁴ Dario Gil, "How to Ensure the U.S.'s Quantum Future," *Scientific American*, August 20, 2020, <https://www.scientificamerican.com/article/how-to-ensure-the-uss-quantum-future/>.

⁵⁵ Lee G. Branstetter, Britta Glennon, and J. Bradford Jensen, "The IT Revolution and the Globalization of R&D," *Innovation Policy and the Economy*, 19 (2019): 1-37, <https://www.journals.uchicago.edu/doi/abs/10.1086/699931>; Britta Glennon, "How Do Restrictions on High-Skilled Immigration Affect Offshoring? Evidence from the H-1B Program" (SSRN, February 21, 2020), <https://ssrn.com/abstract=3547655>; Margaret E. Peters, *Trading Barriers: Immigration and the Remaking of Globalization* (Princeton: Princeton University Press, 2017).

⁵⁶ Tina Huang and Zachary Arnold, "Immigration Policy and the Global Competition for AI Talent" (Center for Security and Emerging Technology, June 2020), <https://cset.georgetown.edu/research/immigration-policy-and-the-global-competition-for-ai-talent/>.

⁵⁷ Zwetsloot et al., "Strengthening the U.S. AI Workforce," 6-7.

⁵⁸ Erica Alini, "2 years in, Trump's immigration squeeze is feeding Canada's tech talent pipeline," *Global News*, April 10, 2019, <https://globalnews.ca/news/5117201/canada-immigration-tech-trump/>.

⁵⁹ Email correspondence with Microsoft Research spokesperson, October 6, 2020.

⁶⁰ John Mannes, "Facebook is the latest tech giant to hunt for AI talent in Canada," *TechCrunch*, September 14, 2017, <https://techcrunch.com/2017/09/14/facebook-is-the-latest-tech-giant-to-hunt-for-ai-talent-in-canada/>.

⁶¹ Branstetter, Glennon, and Jensen, "The IT Revolution and the Globalization of R&D."

⁶² In general, "emerging technologies often [require] a close interaction between R&D and production"; see Börje Johansson and Hans Lööf, "Global Location Patterns of R&D Investments" (Centre for Excellence for studies in Science and Innovation, April 2006), 5, <https://www.researchgate.net/publication/5094381>; A study of deep learning publications and startups suggests that AI developers and AI adopters also benefit from co-locating; see J. Klinger, J. Mateos-Garcia, and K. Stathoulopoulos, "Deep learning, deep change? Mapping the development of the Artificial Intelligence General Purpose Technology," *ArXiv [Cs.CY]* (August 20, 2018), arXiv, 19, <https://arxiv.org/abs/1808.06355>.

⁶³ William C. Hannas, James Mulvenon, and Anna B. Puglisi, *Chinese Industrial Espionage: Technology Acquisition and Military Modernization*, (Routledge, 2013), 65.

⁶⁴ Gert Bruche, "The Emergence of China and India as New Competitors in MNCs' Innovation Networks," *Competition and Change* 13, no. 3 (September 2009): 267–288, https://www.researchgate.net/profile/Gert_Bruche/publication/228714393_The_Emergence_of_China_and_India_as_New_Competitors_in_MNCs'_Innovation_Networks/links/0f317530a2cc960641000000.pdf; Branstetter, Glennon, and Jensen, "The IT Revolution and the Globalization of R&D."

⁶⁵ GlassDoor estimates the average income for a U.S. AI job is \$114,121. "Artificial Intelligence Salaries," Glassdoor, Accessed August 7, 2020, https://www.glassdoor.com/Salaries/artificial-intelligence-salary-SRCH_KO0,23.htm. PayScale reports an average salary for people with AI skills as \$122,127. "Salary for Skill: Artificial Intelligence (AI)," PayScale, Accessed August 7, 2020, [https://www.payscale.com/research/US/Skill=Artificial_Intelligence_\(AI\)/Salary](https://www.payscale.com/research/US/Skill=Artificial_Intelligence_(AI)/Salary). Indeed found that Machine Learning Engineers make an average of \$140,456 per year in the United States. "Artificial Intelligence Salaries in the United States," Indeed, Accessed August 7, 2020, <https://www.indeed.com/salaries/Artificial-Intelligence-Salaries>.

⁶⁶ Brandon Shackelford and Raymond M. Wolfe, National Center for Science and Engineering Statistics, *U.S. Companies Performed \$73 Billion in R&D Outside the United States in 2013* (Washington, DC: National Science Foundation, April 2017), 3, <https://www.nsf.gov/statistics/2017/nsf17317/nsf17317.pdf>.

⁶⁷ Beyond salaries, other cost saving factors include taxes and capital investments (real estate, energy costs, and so forth); see, e.g., Hall, “The Internationalization of R&D.”

⁶⁸ Vijay Govindarajan, Gunjan Bagla, and Avinash Agrawal, “Which U.S. Companies Are Doing the Most R&D in China and India?”, *Harvard Business Review*, March 26, 2015, <https://hbr.org/2015/03/which-u-s-companies-are-doing-the-most-rd-in-china-and-india>.

⁶⁹ Ralph Jennings, “Why Google, IBM & Microsoft Are All Expanding In Taiwan This Year,” *Forbes*, April 4, 2018, <https://www.forbes.com/sites/ralphjennings/2018/04/04/why-google-ibm-and-microsoft-all-decided-to-expand-in-taiwan-this-year/#39bb9fc35c72>.

⁷⁰ Govindarajan, Bagla, and Agrawal, “Which U.S. Companies Are Doing the Most R&D in China and India?”. This cost-benefit calculation may differ for less technical operations, such as sales or certain types of engineering, but our focus is on R&D activities.

⁷¹ Andrew Imbrie, Ryan Fedasiuk, Catherine Aiken, Tarun Chhabra, and Husanjot Chahal, “Agile Alliances: How the United States and Its Allies Can Deliver a Democratic Way of AI” (Center for Security and Emerging Technology, February 2020), <https://cset.georgetown.edu/research/agile-alliances/>; Daniel Kliman, Ben FitzGerald, Kristine Lee, and Joshua Fitt, “Forging an Alliance Innovation Base” (Center for a New American Security, March 2020), <https://www.cnas.org/publications/reports/forging-an-alliance-innovation-base>; Martijn Rasser, Rebecca Arcesati, Shin Oya, Ainikki Riikonen, and Monika Bochert, “Common Code: An Alliance Framework for Democratic Technology Policy” (Center for a New American Security, October 2020), <https://www.cnas.org/publications/reports/common-code>.

⁷² Intel AI Labs: <https://newsroom.intel.com/press-kits/intel-labs/>; NVIDIA AI Labs: <https://www.nvidia.com/en-in/deep-learning-ai/ai-labs-her/>; Uber AI Labs: <https://www.uber.com/newsroom/ailabs/>.

⁷³ The Australian Strategic Policy Institute has already mapped the global presence of China’s largest tech companies, including their R&D labs; see Danielle Cave, Samantha Hoffman, Alex Joske, Fergus Ryan, and Elise Thomas, “Mapping China’s Tech Giants,” Australian Strategic Policy Institute, 2019, <https://www.aspi.org.au/report/mapping-chinas-tech-giants>.

⁷⁴ When it comes to defining “AI,” almost all companies’ labs appear to focus on machine learning-related research.

⁷⁵ Correspondence with Google representatives, September 25, 2020.

⁷⁶ Email correspondence with Microsoft Research representative, October 6, 2020.

⁷⁷ Including affiliates could affect two results: the distribution of staff across companies (Figure 2) and across locations (Figures 3 and 6). For Figure 2, the biggest change when including affiliates is to Amazon’s fraction of total AI staff, which increases from 19 percent to 25 percent of the total. The remaining companies change only by one or two percentage points. For Figure 6, the fraction of staff in any one metropolitan area also varies by no more than one or two percentage points, with the exception of Seattle, which shifts from 21 percent to 14 percent of all staff. The fraction of staff in the United States (Figure 3) changes from 68 percent to 64 percent.

⁷⁸ Zachary Arnold, Rebecca Gelles, and Ilya Rahkovsky, “Identifying AI-Related Companies: A Conceptual Outline and Proof of Concept” (Center for Security and Emerging Technology, July 2020), <https://cset.georgetown.edu/research/identifying-ai-related-companies/>.

⁷⁹ For a general overview of AI R&D metrics, see Raymond Perrault et al., “The AI Index 2019 Annual Report” (Stanford Human-Centered AI Institute, December 2019), https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf. For a broader discussion of metrics to track R&D and innovation, see Bronwyn H. Hall and Adam B. Jaffe, “Measuring Science, Technology, and Innovation: A Review” (May 2012), http://people.brandeis.edu/~ajaffe/Hall-Jaffe%20HJ12_indicators_final.pdf.

⁸⁰ “Business R&D and Innovation Survey (BRDIS),” National Science Foundation, <https://www.nsf.gov/statistics/srvyindustry/#sd&tabs-1>.

⁸¹ Waterloo Artificial Intelligence Institute, “The University of Waterloo’s Artificial Intelligence Institute expands AI for Good partnership with Microsoft Canada,” July 8, 2020, <https://uwaterloo.ca/artificial-intelligence-institute/news/university-waterloos-artificial-intelligence-institute>. Budget figures for Microsoft’s AI for Good program are from John Kahan, “Using AI to Advance the Health of People and Communities Around the World,” Microsoft, January 29, 2020, <https://blogs.microsoft.com/on-the-issues/2020/01/29/ai-for-health-child-mortality/>.

⁸² Lee Branstetter, Guangwei Li, and Francisco Veloso, "The Rise of International Coinvention," in *The Changing Frontier: Rethinking Science and Innovation Policy* (National Bureau of Economic Research, 2015), <https://www.nber.org/chapters/c13028>.

⁸³ IBM, "IBM First in Patents for 23rd Consecutive Year," January 13, 2016, <https://www-03.ibm.com/press/us/en/pressrelease/48785.wss>.

⁸⁴ Nathan Calvin and Jade Leung, "Who Owns Artificial Intelligence? A Preliminary Analysis of Corporate Intellectual Property Strategies and Why They Matter" (Future of Humanity Institute, February 2020), <https://www.fhi.ox.ac.uk/wp-content/uploads/GovAI-working-paper-Who-owns-AI-Apr2020.pdf>.

⁸⁵ Shear, "Facebook Plans To Double Size Of AI Research Unit By 2020."

⁸⁶ "IBM Research | Tokyo," IBM, <http://www.research.ibm.com/labs/tokyo/>.

⁸⁷ Yang, "Microsoft R&D Center in Taiwan Starts Recruiting for AI Research."

⁸⁸ Jay Yagnik, "Google Research India: an AI lab in Bangalore," *Google: The Keyword* (blog), September 19, 2019, <https://blog.google/around-the-globe/google-asia/google-research-india-ai-lab-bangalore/>.