

SEPTEMBER 2020

Patents and Artificial Intelligence: A Primer

CSET Data Brief



AUTHORS

Patrick Thomas
Dewey Murdick

Executive Summary

This data brief informs policymaker audiences who desire to understand how they might use patent data in planning for the quickly advancing impacts of Artificial Intelligence (AI). Such data can provide policymakers with insights into which areas of AI are rapidly developing, which countries are especially active in AI research, and which organizations are responsible for key AI inventions.

In this primer, we report analytic results on worldwide trends in AI patenting and suggest options for how these results might be interpreted and leveraged.

Key findings presented in this primer include:

- There were 10 times as many AI patent applications published worldwide in 2019 as in 2013, most of which have yet to be examined.
- Patent applications increased by 500 percent from 2009 to 2019 within the Chinese patent office—90 percent were domestic applications. The U.S. patent office has seen a 35 percent increase in applications during the same time, 48 percent of which were domestic.
- While the quality of Chinese patents has been repeatedly called into question, there are signs that the situation may be improving.
- Large companies—notably IBM, Microsoft, and Google—dominate AI patenting among U.S. organizations. Meanwhile, Chinese AI patenting is distributed much more broadly across companies (e.g., Ping An, Baidu, Tencent), government organizations (e.g. State Grid), and universities (e.g., Electronic Sci/Tech, Zhejiang, Xidian).
- China focuses AI patenting on Computer Vision, Japan on Control Systems, and Korea on Speech Processing. The United States is more evenly distributed across research fields.

We also place these analytic results in the wider context of patenting and note limitations and considerations with respect to AI patents. Key contextual points found within this primer include:

- Patents are an exchange between the inventor (publicly shared insights) and society (protected rights).
- A patent has a pre-examination “published application” and, if approved, a “granted patent” that confers rights.

- AI patents are strongly dependent on mathematical relationships and algorithms, which are considered abstract ideas under patent law and therefore restrict what can be patented. This concept has and will continue to evolve.¹
- Patent applications have very different meanings depending on where they are filed, which should be considered when comparing innovation trends.

Finally, patent data offers a useful measure of inventive activity across companies, regions, or countries. That said, care should be taken in using AI patent data to support policy decisions, since patenting in AI is growing so quickly. With many recent patent applications yet to be examined, quality and impact are hard to predict.

A Brief History of Patents

At its core, the patent system is designed around an exchange. An inventor receives the right to restrict others from using their invention for a specific period of time. As a result, the inventor obtains the private benefits associated with monopoly-type rights for the life of the patent.² Such benefits may be in the form of licenses from others wishing to use the invention described in the patent, in income generated through litigation against those who use (“infringe”) the patent without permission, or in injunctions preventing infringers from using the patented invention.

In return for these private benefits, the inventor must describe the invention in sufficient detail that others skilled in the profession can understand and potentially further develop the idea. The economy thus receives the benefit of technology advancing at a faster pace than if everybody kept their innovations secret. Hence, the private benefit to the inventor also generates a public benefit for society.

Patents have a long history. The first patent was granted in 1421 in Florence, Italy to an inventor of a barge fitted with hoisting gear to lift marble. England first started issuing patents in the early 17th century, and the United States granted its first patent in 1790, for a process to make potash, an ingredient used in fertilizers. The history of patents in Asia is shorter, in some cases strikingly so. While the nineteenth century saw the establishment of patent systems in India (1856) and Japan (1885), patent offices were not established until recently in South Korea (1977) and—perhaps most importantly in the context of AI patenting—China (1980).

For most of their history, patents have been largely directed to items one could see and touch (e.g., Edison’s light bulbs, Hargreaves’ spinning jenny, Bell’s telephones), and processes for manufacturing and transporting these items. This reflected the economic structure of many societies following the Industrial Revolution, with inventors filing patents to cover their innovations in the most profitable sectors. The digital revolution has upended those economic structures in many advanced societies. Many of today’s most lucrative innovations cannot be seen or touched, but instead represent intangible concepts embedded in data processing and software. In turn, the number of patents filed for intangible inventions, such as AI, has grown rapidly. This shift has had major implications for patent systems worldwide, as they work out how to address new types of innovations.

Can AI Be Patented? ³

While the concepts behind the patent system are relatively simple, the rules associated with obtaining a patent are much more complicated, and vary widely across the world. That said, standard elements exist among all major patent systems. The patenting process starts with an applicant filing a patent application. This application describes their invention, and—in the claims section—delineates the specific scope of the invention being claimed. In most patent systems, patent applications are confidential for 18 months—after which the “published application” is disclosed.

A patent examiner analyzes the patent application to determine whether it meets all of the applicable requirements. This examination is typically not a one-time event, but rather an extended negotiation between the applicant and examiner referred to as “patent prosecution.” The negotiation either results in a set of claims acceptable to the patent examiner or it does not. If it does, the patent application is allowed, resulting in a granted patent. Hence, two distinct documents may be associated with a given patent—a pre-examination “published application” and (if deemed allowable by the examiner) a “granted patent.”

Patent examination is founded on a core group of principles across all national and international patent systems—that, to be patentable, an idea should be new, have some practical use, be non-obvious, and not be naturally occurring. To use the U.S. system as an example, applicants must satisfy the following requirements to obtain a granted patent, each associated with a particular statute in U.S. law:

1. **Patentable subject matter** (Title 35 U.S.C. § 101): This statute covers two concepts. The first is “utility”—i.e., the invention must have some practical use. The second is that the invention must describe subject matter deemed to be patentable, namely “any new and useful process, machine, manufacture, or composition of matter.” The guiding principle here is that “anything under the sun that is made by man is patentable.”⁴
2. **Novelty** (Title 35 U.S.C. § 102): The invention must be novel (i.e., it must not have been known publicly prior to the application for the patent). In simple terms, one cannot get a patent for an invention that is not new.

3. **Non-obviousness** (Title 35 U.S.C. § 103): The invention must not be an improvement to the existing state of the art that would have been obvious to practitioners of “normal skill” in the field. This disallows patents for inventions that are minor alterations to existing technologies, or that incorporate obvious extra elements. Non-obviousness often results in extensive negotiation between the applicant and examiner during patent prosecution.⁵

The “patentable subject matter” requirement has proven particularly challenging for patent offices and courts deciding on intangible inventions associated with the digital revolution.⁶ The “made by man” requirement means various subject matters are not patentable, notably laws of nature, abstract ideas, and natural phenomena.⁷ AI runs particularly close to this subject matter, especially since mathematical relationships and algorithms are considered abstract ideas under patent law.

Guidance related to the patentability of intangible concepts has evolved in the United States over the past two decades in response to rulings from the court system. Evolution is commonplace with emerging technologies, as the patent system reacts to new developments that were not envisioned when patent law was established. In the case of intangible concepts, following the 1998 “State Street” case,⁸ many such concepts became patentable, provided they were implemented on a computer. This resulted in a spike in patent applications, especially related to business methods. Subsequent rulings—notably in the 2014 “Alice” case⁹—have stated that taking an abstract idea, such as an algorithm or a business method, and implementing it on a generic computer is not sufficient to be patentable.

Under current guidance,¹⁰ patent examiners must use a multi-step, multi-prong process to determine whether an AI invention is patentable. They must first decide whether the invention is directed to an ineligible concept, notably an abstract idea in the case of AI. If it is, they must then use a multi-prong step to address a number of other factors. These include whether the invention is integrated into a practical application (e.g., improving the functioning of a machine or aiding the treatment of a disease or condition). Another prong considers whether the invention contains an additional element or is directed to a particular application that makes the invention more than a claim on the abstract idea itself (with implementation on a generic computer no longer being sufficient to satisfy this step).

In simple terms, an applicant for a patent related to AI must demonstrate that the invention amounts to more than simply a claim on the underlying AI

algorithm itself. If the applicant can clear the various hurdles to achieve this, a patent is granted, and offers the same basic protections as all other patents from the same issuing authority—in our example above, the U.S. Patent & Trademark Office (USPTO). In this way, in legal terms, a patent for a neural network application becomes like any other patent, whether for a wireless router, a pharmaceutical composition, or indeed a wheelbarrow.

How Does AI Patent Protection Work?

Earlier, we introduced the idea of inventors receiving protection for their innovations via the patent system. There are three important general characteristics of patent protection:

1. **It is limited geographically:** The coverage of a patent is limited to the jurisdiction of its issuing authority. For example, a patent granted by the USPTO (a “U.S. Patent”) provides protection only within the United States. If an organization intends to protect an invention in multiple countries, it must file patents in each of those countries’ systems. As a result, patent applications filed in a particular country’s system are not synonymous with inventions made in that country. Indeed, according to USPTO, roughly half of all patent applications it receives are from overseas (based on location of first inventor).
2. **It is limited temporally:** Patent protection is not endless, but has a time limit. The private benefits accruing to the inventor would otherwise greatly outweigh the public benefits enjoyed by society. Currently in the United States, the term for utility (i.e., “invention”) patents is typically 20 years from the date the application is filed.

In contrast, in China (and in many other countries, including Japan, Korea, and Germany), two different types of patent protection exist for inventions.¹¹ The first is an “invention patent,” which equates to a U.S. utility patent and has a lifetime of 20 years. The second is a “utility model.” It has a lifetime of only 10 years and undergoes a less rigorous examination than an invention patent prior to being granted. Utility models are generally less expensive to file than invention patents, and have a greater chance of being granted. Often, they are used to protect lower-profile inventions or inventions for which long-term protection is not seen as necessary.

With respect to AI, one important feature of utility models in China is that they can only be used to protect innovations with a physical shape and structure. This requirement excludes methods and processes, as well as intangible items such as software. Hence, while utility models outnumber invention patents in the Chinese patent system in general, they represent only around 3 percent of the total in AI.

3. **It is limited to granted patents:** Filing a patent application alone does not generally accord an inventor the private benefits associated with a patent. Otherwise, inventors could apply for extremely broad patents covering wide swathes of the economy. Instead, as noted above, patent applications go through an examination process. Only if applications succeed in this process are patents granted and become enforceable against others.

The rules for examining and enforcing patents vary across patent systems. In the United States, all patent applications are examined (unless the applicant withdraws them prior to examination) in an automatically initiated process. Currently, the mean pendency (i.e., the time from application to grant) for U.S. patent applications is just under two years—a reduction in recent years as the USPTO has made concerted efforts to shorten pendency periods (e.g., it was over 27 months in 2014).

The examination system is different in China, as well as in Japan and South Korea. Specifically, for invention patents, applicants have three years to request examination of their application. Without a request, the application is abandoned after that three-year period. Patent application pendency in China is generally shorter than in the United States; the Chinese Patent Office (CNIPA) announced a target of 16 months by 2022. Hence, the time from application to grant in China ranges from less than two years (if the applicant requests examination immediately) to more than four years (if the applicant takes the full three years to request examination).

How Are AI Patents Defined and Characterized?

To study trends in AI patenting, CSET and 1790 Analytics created a [database](#) containing all AI granted patents and published applications (collectively “patent documents”) between January 2000 and March 2020

in all patent systems worldwide.¹² We identified AI patents using a combination of keywords and patent classifications, notably the Cooperative Patent Classification and International Patent Classification.¹³ Our AI patent database contains 287,532 patent documents (230,855 published applications and 56,677 granted patents). These patent documents are grouped into “patent families,” with each family containing all patent documents associated with the same original invention (e.g., if patent applications are filed in multiple countries).¹⁴ There are a total of 155,770 patent families in the database.

The patent families in the AI patent database are categorized along three dimensions:¹⁵

- AI Techniques: How does the invention work? (e.g., machine learning, logic models, fuzzy logic)
- Functional Applications: What does the invention do? (e.g., speech processing, computer vision, control systems)
- Application Fields: Where can the invention be used? (e.g., life sciences, transportation, energy management)

Within each dimension, a patent family may be included in more than one category. For example, a family describing an AI application for encrypting communications may be included in both the Security and Telecommunications categories within the Application Fields dimension. A patent family may also be classified along more than one dimension. For example, a family describing a system for controlling an autonomous vehicle using a machine learning-based speech recognition tool will have entries in all three dimensions (i.e., AI Technique: Machine Learning; Functional Application: Speech Processing; Application Field: Transportation).

The database includes extensive details on all AI granted patents and published applications, including: application and publication dates; priority and patent family information; patent classifications; assignees (many of which are normalized and account for subsidiary names and mergers and acquisitions, resulting in accurate patent portfolios for organizations); titles; abstracts; and technology category information for each of the three dimensions outlined above.

What Are the Analytic Limits of AI Patent Data?

Patents have long been used to measure inventive activity, whether of companies, regions, or countries. One advantage of employing patent data is that it represents an output measure from research, in contrast to input measures such as R&D expenditures. Patents are also inherently technological, unlike information extracted from other sources. For example, a patent for an AI-controlled robot must by definition describe the technology associated with this robot, while other information sources may focus on its commercial, ethical, and social implications.¹⁶

While patents are a useful measure of inventive activity, employing them to measure competitiveness in technologies such as AI requires particular care. Many patent-based analyses focus on technologies and patent systems that are relatively stable, allowing researchers to evaluate trends within this “steady-state” context. For example, if Japanese automakers file increasing numbers of transmission patent applications, while German automakers’ patenting trends toward combustion, it is possible to determine where the two countries are focusing their research efforts.

“Steady-state” is the perhaps the last term that could be associated with the Chinese patent system or with the contours of AI technology in general. Both are changing quickly, which has major implications for using patent data to assess AI innovativeness—both over time and across countries. In particular, it raises the question of how to assess the dramatic recent increase in Chinese AI patent applications within this rapidly changing landscape.

According to CNIPA, the number of invention patent applications it received increased by more than 500 percent between 2009 and 2019, from 241,000 to 1.4 million (although, interestingly, there was a 9 percent decrease from 2018 to 2019). In comparison, the number of patent applications at the USPTO increased by only 35 percent (from 456,000 to 621,000) over the same time period. Hence, while in 2009 U.S. patent applications outnumbered Chinese applications by almost two-to-one, by 2019, the ratio had completely reversed.¹⁷ Most of the Chinese patenting increase can be attributed to applications filed by domestic applicants. Out of the 1.4 million CNIPA applications in 2019, domestic sources filed almost 90 percent (compared to 48 percent of USPTO applications).

Numerous writers have pondered whether this rapid growth in Chinese patenting in fact reflects increased inventive activity, or results from other

factors. Particular prominence is given to the government subsidies and tax breaks available to Chinese companies filing patent applications. These subsidies form part of the “Made in China 2025” initiative, which aims to move the Chinese economy further up the value chain, from a low-cost manufacturer of products designed elsewhere to an independent source of high-value products and services. AI plays a critical role within this initiative.

In order to qualify for these subsidies, some suggest that Chinese companies have prioritized quantity over quality in their patent applications, resulting in a flood of low-quality Chinese patent applications (notably, the U.S. patent office has not been immune from accusations related to patent quality). Two statistics have been used extensively to support this assertion. The first is the “success rate” for Chinese applications being granted. For example, in 2017, only 26 percent of Chinese patent applications from domestic residents were granted, versus 68 percent for foreign residents (whereas for U.S. patents, the success rate for domestic and foreign applicants is roughly equal, at just over 50 percent).¹⁸ Lack of overseas patenting by Chinese inventors has also been used to suggest that the poor quality of Chinese patent applications means they may not withstand examination by other patent systems. For example, in 2016, only 4 percent of Chinese patent applications were filed overseas, versus over 40 percent for U.S. applications.¹⁹

Statistics such as these have been used to paint Chinese patent applications as largely worthless. However, evidence also suggests that—to the extent they exist—any problems related to Chinese patent quality may be in the process of being addressed. In 2019, CNIPA examined 1.02 million invention patents and granted 453,000, a “success rate” of 45 percent.²⁰ This percentage is in line with more established patent systems, including in the United States (although it does not address applications for which no examination has been requested or potential differences in patent laws across systems). Additionally, while the success rate at CNIPA is not reported separately for domestic versus foreign applicants, given an overall success rate of 45 percent, the success rate of domestic applications is likely to be much higher than the 26 percent reported for 2017.

Chinese inventors are also becoming more active in filing patent applications beyond their domestic system. In 2019, China led in patent applications under the Patent Cooperation Treaty (PCT) system of the World Intellectual Property Organization (WIPO) for the first time.²¹ This broke the United States’ long run as the leader in this system since its establishment in 1978. PCT applications streamline the filing of patents internationally, so China’s

increase in filings (from 26,000 in 2014 to 59,000 in 2019) shows a growing commitment of Chinese applicants to protect their innovations abroad.²² The number of foreign patent filings by Chinese applicants in individual countries also more than doubled between 2013 and 2018, from 29,000 to 66,000. At the same time, foreign patent filings by U.S. applicants rose less than 10 percent, from 213,000 to 230,000.


This recent growth in overseas patenting by Chinese applicants can also be seen in the AI patent data. In our AI patent database, we identified 2,894 PCT applications that have a Chinese priority application (i.e., where the invention was filed originally in China). More than 60 percent of these original Chinese applications are from the period since 2017. Also, there are 5,237 non-Chinese AI patent applications (excluding PCT applications) that resulted from an original Chinese priority application, more than 80 percent of which have been published since 2017.²³ These statistics suggest that recent years have not only seen an increase in domestic Chinese AI patenting, but also an extension of this patenting beyond the domestic patent system.²⁴

These statistics suggest the Chinese patent system is developing characteristics of more established patent systems, though progress remains; for example, the proportion of PCT applications versus domestic applications is still much lower in China than in the United States. As such, while concerns about Chinese patent quality are well-founded, Chinese patent applications should not be dismissed. In particular, the next few years will be instructive, especially in AI, as the rush of recent Chinese patent applications are either granted or abandoned—a significant signal of the value to be attached to Chinese AI published applications.

What Meaning Can Be Extracted from AI Patent Data?

In order to extract useful insights from AI patent data, one must present the findings carefully and consider the concerns discussed above. For example, Figure 1 shows the number of AI published applications and granted patents worldwide each year between 2000 and 2019.²⁵


Figure 1 – Number of AI Published Applications and Granted Patents by Year



This figure reveals the sharp increase in AI patent activity in recent years. There were 10 times as many AI patent applications published in 2019 as in 2013. The same time period saw an almost four-fold increase in granted AI patents. Presenting granted patents and published patent applications separately provides further insights into trends in AI patenting. In particular, it emphasizes that a great deal of AI patent activity is very recent, as reflected by the sharp increase in published applications in 2018 and 2019. Meanwhile, the increase in granted AI patents in 2019 suggests that the applications filed in the early part of the recent spike are starting to work their way through patent prosecution. Revisiting a chart such as this in the coming years will provide an interesting benchmark for assessing how many of the tens of thousands of recent AI patent applications have resulted in granted patents.

Patent data can also be used to identify organizations with extensive AI research interests. Figure 2 shows the organizations with the largest number of AI patent families.


Figure 2 – Number of Patent Families (Jan 2000–Mar 2020) Owned by Leading AI Organizations (Red = China; Blue = United States; Green = Other)



Beyond the individual names in Figure 2, a number of interesting insights can be extracted from this chart. The first is that the largest AI granted patent portfolios are associated with major U.S. technology companies—namely, IBM, Microsoft, and Google. However, beyond these three companies, Chinese organizations dominate. These include companies (Ping An, Baidu, Tencent), government organizations (State Grid), and universities (Electronic Sci/Tech, Zhejiang, Xidian), reflecting the breadth of AI research in China. Notably, the patent portfolios of these Chinese organizations currently consist almost entirely of published applications, rather than granted patents. Many of these Chinese published applications are very recent. Hence, as with the trends in Figure 1, it will be interesting to revisit this chart in the coming years to determine how many of the published applications associated with the different Chinese organizations have become granted patents.


Patent data can also provide insights at a national level. Figure 3 shows the number of AI patent families by priority country.²⁶ This figure reflects the dominance of China and the United States in terms of AI patent activity. While having the largest number of families overall, China has a much lower percentage of families containing a granted patent than the United States. This may again result from the different age profiles of the two portfolios, with most of China’s AI patent families filed in the last three years, whereas the United States has a much longer history of patenting in AI.²⁷

Figure 3 – Number of AI Patent Families by Priority Country




Evaluating patents at a more granular level is also instructive. For example, Figure 4 shows the distribution of patent families across AI functional applications for leading countries.²⁸ This figure reveals the relative focus of China on Computer Vision, Japan on Control Systems, and South Korea on Speech Processing.

Figure 4 – Distribution of Patent Families across Functional Applications by Priority Country



Counting AI patents in different ways allows for analyses of trends in AI patenting across countries, organizations, and technology categories. Moving beyond patent counts, advanced metrics can provide insights into the relative strength of different portfolios. For example, Figure 5 shows the Citation Index values for leading organizations.²⁹ It suggests Google and Amazon have relatively high-impact AI patents, whereas IBM, Sony, and Siemens fare less well.³⁰

Figure 5 – Citation Index Values for Leading Organizations (based on AI granted U.S. Patents)



Acknowledgments

For feedback, reviews, and other assistance, we would like to thank Zach Arnold, Jill Crisman, Melissa Flagg, Mark Fleming, Rebecca Gelles, Saif Khan, Igor Mikolic-Torreira, Santiago Mutis, Ilya Rahkovsky, Richard Silbergitt, and Alexandra Vreeman.

Patrick Thomas is Co-Founder and Director of Analytics at 1790 Analytics.

Dewey Murdick is the Director of Data Science at the Center for Security and Emerging Technology.



© 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/>.

Document Identifier: doi:10.51593/20200038

Endnotes

¹ United States Patent and Trademark Office, Commerce. “2019 Revised Patent Subject Matter Eligibility Guidance.” Federal Register 84, no. 50 (January 7, 2019): 2018-28282. <https://www.federalregister.gov/documents/2019/01/07/2018-28282/2019-revised-patent-subject-matter-eligibility-guidance>.

² We use the term “monopoly-type rights” rather than simply “monopoly rights” as patents give their owners the right to exclude others, not necessarily to use the invention themselves, since they still risk infringing other patents.

³ Separate from the issue of whether AI *can* be patented is a broader philosophical question of whether AI *should* be patented, especially given the speed at which the technology develops. This broader question is beyond the scope of this primer, which focuses primarily on the empirical question of how AI patent data may be employed by policymakers, rather than the philosophical question of whether such data should exist at all.

⁴ *Diamond v. Chakrabarty*, 447 U.S. 303 (1980).

⁵ Non-obviousness is the most common reason for first final rejections of patent applications. Despite its ominous name, a first final rejection is not necessarily the end of patent prosecution. The applicant may amend and re-submit their application in response to a final rejection, but must pay additional fees if any changes are made to the claims. For more on this process, see <https://www.ipwatchdog.com/2016/10/03/103-rejections-common-respond/id=73214/>.

⁶ Note that non-obviousness also poses significant challenges in AI, and its application to different types of AI inventions could lead to certain areas of AI being over- or under-represented in the patent data.

⁷ Interestingly, patent offices are now wrestling with a new question of whether inventions made by intelligent machines (rather than “made by man”) should be patentable, as in the case of a recent patent application that listed the inventor as DABUS, a “creativity machine”: <http://artificialinventor.com/dabus/>.

⁸ *State Street Bank and Trust Company v. Signature Financial Group, Inc.*, 149 F.3d 1368 (Fed. Cir. 1998).

⁹ *Alice Corp. v. CLS Bank International*, 573 U.S. 208 (2014).

¹⁰ United States Patent and Trademark Office, Commerce. “2019 Revised Patent Subject Matter Eligibility Guidance.” Federal Register 84, no. 50 (January 7, 2019): 2018-28282. <https://www.federalregister.gov/documents/2019/01/07/2018-28282/2019-revised-patent-subject-matter-eligibility-guidance>.

¹¹ Patents related to inventions dominate all major patent systems (e.g., 93 percent of U.S. applications in 2019 were for utility patents). They are generally what people think of when referring to patents, especially in a technology area such as AI. There are also other types of

patents in different systems. For example, in the United States, there are design patents (for aesthetic product designs) and plant patents (for reproducible plants, seeds, tubers, etc.).

¹² A detailed description of this classification along with an index for patent documents filed around the world can be found on GitHub: <https://github.com/georgetown-cset/1790-ai-patent-data>.

¹³ The keywords and patent classifications used to identify AI patents are adapted from an earlier study of AI patenting carried out by the World Intellectual Property Organization (WIPO), available at https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf. In our analysis, we used a more tightly defined set of classifications and keywords, resulting in a somewhat smaller and more focused final patent set.

¹⁴ Patents are grouped into families based on sharing the same “priority document,” which is the initial application in the family.

¹⁵ The three dimensions are taken from the WIPO study, with some categories added to the dimensions to reflect recent technological developments (e.g., Semiconductors, Nanotechnology). The patent classifications and keywords used to allocate patents to categories were developed for this analysis. In turn, the technology categories in both this analysis and the earlier the WIPO study build upon the Computing Classification System (CCS) developed by the Association for Computing Machinery (ACM).

¹⁶ See Dubbert et al. “Using Intellectual Property Data to Measure Cross-border Knowledge Flows”, USPTO Economic Working Paper No. 2019-02, March 2019, for a useful overview of the different ways in which patent data can be used.

¹⁷ In addition to China’s 1.4 million invention patent applications, there were almost two million utility model applications in China in 2019.

¹⁸ Peter Finnie, “Why China’s impressive patent rates don’t tell the whole story.” New Statesman Tech, February 11, 2019, <https://tech.newstatesman.com/guest-opinion/china-patent-rates>.

¹⁹ Ana Maria Santacreu and Heting Zhu, “What Does China’s Rise in Patents Mean? A Look at Quality vs. Quantity,” Economic Synopses, No. 14, 2018. <https://doi.org/10.20955/es.2018.14>.

²⁰ Aaron Wininger. “Chinese Patent Filings Drop 9% in 2019.” China IP Law Update, January 14, 2020, <https://www.chinaiplawupdate.com/2020/01/chinese-patent-filings-drop-9-in-2019/>.

²¹ Based on location of the first-named applicant.

²² PCT applications do not themselves mature directly into patents, but give the applicant more time (30 months) to decide the countries in which to file an application for their invention. Given that many Chinese PCT applications are still within this 30-month window, it remains an open question as to how many will result in foreign patent applications.

²³ In comparison, a total of 13,487 AI PCT applications have a U.S. priority application, with 37 percent of these U.S. applications filed since 2017. There are also 34,713 non-U.S. AI patent applications (excluding PCT applications) that resulted from an original U.S. priority application, half of which have been published since 2017.

²⁴ Note that all of these statistics may in fact understate the recent growth in AI patenting. Given that patent applications are generally confidential until being published after 18 months, the data will not include many recent applications that have yet to be made public.

²⁵ This figure counts all published applications and granted patents separately in order to show their relative trends. As a result, there may be some double-counting of published applications and granted patents from the same family.

²⁶ Allocating patents to countries is not straightforward. Large companies may have inventor teams from multiple countries, plus they may file patents via holding companies (e.g., Alibaba has some patents filed under an assignee in the Cayman Islands). Extending beyond individual patents to families further complicates this, since inventor teams may alter somewhat, and different assignees may be used for different documents in a family. Also, from a practical perspective, assignee and inventor location information is not always available immediately for all patent systems. There is also a philosophical question, in terms of which country should get credit for an invention, for example, if IBM files a patent based on work done in its labs in China (and if this should be treated differently than a patent invented and assigned to a Chinese company). To simplify the analysis, our approach is to allocate patent families to countries based on the priority country (i.e., where the first patent application in the family was filed; in cases where the first application is a PCT application, the priority country is taken from the country code in the PCT number). We ran a test of this approach, using a sample of 7,335 machine learning patent applications published in 2019. This test revealed an 85 percent match between priority country and assignee country, and an 81 percent match between priority country and first inventor country. These percentages are comparable to the match between assignee country and first inventor country (81.5 percent), reflecting the lack of a perfect answer in allocating patents to countries.

²⁷ In the past, analysts sometimes gave greater weight to “triadic” patents—i.e., those filed in the United States, Europe, and Japan. China could potentially be added to this list, depending on how its patent system develops in the future.

²⁸ For presentation purposes, this figure does not differentiate between families with and without granted patents. Also, the purpose of the figure is to show the relative technological focus of different countries, not to benchmark them in terms of their strength in different areas.

²⁹ The Citation Index is a normalized impact metric based on forward citations. It has an expected value of one. Values above one show a patent portfolio cited more than expected; below one shows fewer citations than expected.

³⁰ The metric is based on granted U.S. AI patents. The absence of Chinese organizations is due to their relative lack of such patents, but these organizations may qualify in the future as their portfolios of granted U.S. AI patents grow.