

LBS Research Online

[M G Jacobides](#), S Brusoni and F Candelon

The evolutionary dynamics of the artificial intelligence ecosystem

Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/1789/>

[Jacobides, M G](#), Brusoni, S and Candelon, F

(2021)

The evolutionary dynamics of the artificial intelligence ecosystem.

Strategy Science, 6 (4). pp. 412-435. ISSN 2333-2050

DOI: <https://doi.org/10.1287/stsc.2021.0148>

INFORMS (Institute for Operations Research and Management Sciences)

<https://pubsonline.informs.org/doi/abs/10.1287/sts...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.

1. Introduction

In recent years, the emergence of Artificial Intelligence (AI) has generated excitement and concern in equal measure. Such mixed emotions about the potential impact of a form of human-made yet non-human cognition have been reverberating since the 1950s. Yet, the context for today's discussion is dramatically different, because it unfolds in parallel to the actual application of AI-based technologies to everyday life. AI is no longer confined to the lab, or specialized applications in some esoteric scientific field, or a super-computer challenging a chess grandmaster. A whole range of AI-enabled products and services are on the market right now, from search engines to face recognition, call-center chatboxes and bots to medical diagnosis and autonomous driving—and more will soon emerge. Hence, the conversation has shifted from a highbrow debate about the nature of intelligence and humanity to a practical discussion of business models, regulation, ethics, data property rights, reskilling, and the impact on employment structures. However, the intellectual conflict over the nature of intelligence still persists—and with good reason.

Today, AI is a pressing priority. The World Economic Forum, in its 2018 *The Future of Jobs* report, identified AI as the core of a cluster of related technologies (including high-speed mobile internet, Cloud computing, and big data analytics) that “are set to dominate the 2018–2022 period as drivers positively affecting business growth.” Predicting a rapid and accelerating pace of adoption, the report stressed the implications for employment trends and firms' development strategies. AI has also generated considerable practical excitement for firms; Iansiti and Lakhani (2020) posit that “markets are being reshaped by a new kind of firm—one in which artificial intelligence runs the show,” with different underlying economics and organizing principles, which they explicate. AI is also treated as a priority by entire countries. As Babina et al. (2020) note, the US government is looking to double its non-defense research and development (R&D) budget for AI (Executive Office of the President, 2019); the European Union called for a \$24 billion investment in AI research by 2030 (European Commission, 2020); and China is aiming to invest \$150 billion in its domestic AI market by 2030 (Mou, 2019). Furman and Theodoridis (2020) show that AI can also make researchers more productive. That said, some research is more circumspect. Brynjolfsson, Rock, and Syverson (2017) point out that aggregate productivity growth has actually slowed down in recent times (pre-pandemic), despite the increasing availability of so-called “transformative technologies.”

One of the challenges of existing research is that it either focuses on *aggregates*, or on specific *applications* enabled by AI. For example, following a longstanding tradition, economic research on AI (e.g., Aghion et al., 2019) explores neither who produces AI, nor

who consumes it. Instead, its focus (see Aggarwal et al., 2019 for an excellent review) is AI's aggregate impact on productivity and the jobs market (Tambe et al. 2019; Furman & Seamans, 2020). It also considers whether AI is broad enough to qualify as a "General Purpose Technology" (GPT), per Bresnahan & Trajtenberg (1995), which also suggests that society would be better off finding ways to support its development—since, the theory goes, the benefits of AI are too diffuse to be privately funded. Conversely, research in management tends to focus on specific applications of AI (e.g. health care applications (Garbuio and Lind, 2019; Allen et al. 2019), media industries (Chan-Olmsted, 2019), academic research (Furman and Teodoris, 2021) etc.) or specific managerial activities (e.g. business model innovation (Burstroem et al. 2021), organizational decision-making (Sherstha et al. 2019), and marketing (Kumar et al. 2019).

Both these approaches are useful, yet neither provides a map of the key actors in the world of AI and their business models, in the key geographies. We think that an evolutionary approach, which focuses on opportunities to generate and appropriate value, and shows how firms generate, support, and apply AI can help better describe, understand, and prescribe. As management and strategy scholars, we wish to look beyond the broad categorization of AI "as just another chapter in the 200-year story of automation" (see, e.g., Aghion et al, 2019: 238). We focus on the emerging division of labor between different types of firms that engage in AI, moving beyond secondary data reports (see Simon, 2019).

Our aim is not merely to map AI actors, but also to consider the meso-level, mid-term evolutionary processes that support the development and application of AI, which help us understand what AI *is*, who *produces* it, and who *benefits* from it. To that end, we set out an evolutionary account of the *AI innovation and production system*—that is, the network of interconnected organizations and institutions that is enabling its rise. We build on past and ongoing work on national and sectoral systems of innovation (e.g., Lundvall, 1992; Malerba, 2004), i.e. a set of functionally connected, yet heterogeneous actors (e.g., firms, communities, networks) and institutions (e.g., governments, public and private research labs) that operate on the basis of common bodies of knowledge and sets of technologies (in our case, AI-related). Following Jacobides and Winter (2005,2012), and Jacobides et al (2006) we consider the nature of the "industry architectures" that emerge in AI and look at how they differ in terms of the key players. We then ascend from these "roots" of evolutionary theory to more recent "branches" of research on triple-helix, business and innovation ecosystems, and on digital platforms (Etzkowitz & Leydesdorff, 1995; Adner, 2017, Cusumano & Gawer, 2014).

The two broad questions we ask are "What is the nature of AI, and of the actors engaged in its production and consumption?" and "How does AI affect the evolutionary dynamics of firms

and industries, in the key national settings?” To our knowledge, our paper provides the first comprehensive systematic analysis of the activities involved in the production, enablement, and consumption of AI, and a comprehensive overview of the main players and their business models, drawing on direct evidence (Simon, 2019). Our overview raises a number of issues.

First, we question the view that AI, as a GPT, should receive public funds and emphasize that AI is very unevenly distributed. We showcase the remarkable and growing concentration in AI, and the fact that many Big Tech firms span all the way from infrastructure to applications, leading much of the relevant scientific advantage, too. We also consider some strategy questions, such as whether AI leads to firms migrating to different sectors. In all, we argue that to answer policy and strategy questions, we need to understand how the shifting economics of AI shape its evolution and development, and study how firms’ strategies can shape future technologies and their downstream application— which is what this paper offers.

The remainder of the paper unfolds as follows. Section 2 looks at AI as a technical system, while Section 3 examines *who* undertakes each activity within it. Section 4 zooms in on the dynamics of AI’s downstream application to understand the actual dynamics of AI provision. Section 5 highlights the differences between AI’s evolutionary trajectories in the U.S., China, and the EU. Finally, Section 6 summarizes our theoretical contribution, and Section 7 considers how our approach differs from, and complements, existing ones.

2. AI (and ML in particular) as a technical system

Herbert Simon (1970), in his landmark book on the “Sciences of the Artificial,” argues that humans have been able to advance largely by creating “artificial” worlds (contrasted with the “natural” worlds they inhabit) by engineering structure and creating systems whose objective is to adapt. AI is one such system but what exactly *is* its “artificial” structure? How can we comprehend AI as a technical system (Hughes, 1993; Rosenberg, 1982)? What steps are involved in its production and consumption? As Cockburn et al. (2019) note, we can categorize AI into three areas: symbolic systems, robotics, and machine learning (ML). We focus our attention on the last of these, where most progress has recently taken place.

As several (excellent) conceptual reviews already exist (e.g. Marcus, 2018), we will not attempt one. Nor will we propose another typology of how AI affects firms, since major consultancies have made headway here (Gerbert et al, 2017 [BCG]; Bughin et al, 2019 [McKinsey & Co/MGI]; Sudarshan et al, 2018 [Deloitte]; Herweijer et al, 2018 [PWC/WEF]; Ransbotham et al, 2020 [BCG / MIT]). Instead, our focus is to understand how AI (and especially ML) is produced and consumed. To do so, we move beyond the reliance on

secondary sources (Simon, 2019), and draw on a number of projects undertaken by the authors and their institutions, including a large-scale survey and 37 semi-structured interviews with senior executives in a number of the key actors in AI, as well as with executives from GAMMA, BCG's specialist arm on digital transformation and AI, where 30 projects were reviewed. Further details are provided in the Online Appendix.

2.1. What steps allow us to enable, produce, and consume AI/ML?

Turning to our setting, Figure 1 below illustrates the steps and effort flow involved in the development and operation process for ML.¹ The figure illustrates the different activities that need to be undertaken, from data sourcing and integration, to model creation, training and continuous monitoring and improvement after release.

Insert Figure 1 around here

This technical decomposition of AI/ML development, however, is a fairly granular picture, and tells us little about the typology of players interacting at the level of the AI ecosystem. Enabling and using AI is more than just choosing the right type of algorithm and developing it. It involves other components, such as hardware, data management, AI platforms, and AI applications. We provide a rough illustrative version in Figure 2 below, drawn from our engagement with AI specialists, analysts, and consultants and the specialist literature. First, AI technology is underpinned by *enablers*, which include physical infrastructure (e.g., chip technology) and data management and processing. Second, AI enablement supports the AI development environment, which encompasses platform technologies (e.g., AWS or Google TensorFlow) or other visualization software (e.g., Facets, TensorWatch, Tableau). Third, AI use cases developed in these environments can be deployed in conjunction with industry-specific applications to support businesses in optimal resource allocation or personalization (Candelon et al. 2020). We consider AI consumption to be the use of analytical solutions in an industry application, thus turning the latent possibilities of AI into a specific output.

Include Figure 2 about here

What do we learn from Figure 2? First, AI consumption and production are intricately connected in some key segments, particularly ML: AI consumption (i.e., using an algorithm) can provide the data to calibrate it, and this leads to a positive feedback loop. This makes AI unique among GPTs: no steam-engine or electric-motor output would endogenously improve

¹ This figure is consistent with recent work in the computer science literature (see e.g., He et al. 2020: Figure 1).

by being put into use, setting aside learning-curve and application improvements.

Second, AI enablement, production, and consumption lead to significant demand for computing power, which is provided primarily by Cloud computing companies. BCG estimates that the AI-generated demand for Cloud servers and platforms amounted to \$5.7 billion in 2020, and that sum is expected to grow by 43% per year in the next three years.² Therefore, Cloud providers will be keen to resolve any downstream bottlenecks— analogous to what Ethiraj (2007) found in an earlier period of the computer sector evolution, when firms in one part of the value chain would help innovation in another part, that was holding them back (also see Baldwin, 2018).

Third, AI depends on good-quality data (Economist, 2017), and firms that own, or can access, this vital resource are more likely to engage in AI. As sensor technology improves and the Internet of Things (IoT) becomes more prevalent, the volume of data will only increase; every change in the physical world will be reflected in the digital world. Given that data is non-rivalrous, the same information can theoretically be used by many applications simultaneously. Yet the way data is accessed, and the ability to draw on appropriately structured datasets, is becoming a source of competitive advantage, given its ability to leverage AI. Future regulations might set data free in order to enable competition—but until then, access to data is a key strategic factor and enabler of AI.³

Finally, communities of programmers complement the key actors throughout the AI space. For example, AI communities such as those on GitHub (acquired by Microsoft in 2018) and Kaggle provide an online space where developers can access and contribute to myriad datasets, algorithms, and models, and advance their AI knowledge through online courses.⁴ Hence, they serve as a flexible, dynamic platform to spur AI innovation and commercialization.

2.2. Understanding the role of AI libraries

² Note that this is only a small part of the Cloud demand, but it is one that is critical to the success of Cloud providers' clients (e.g. think of the importance of Netflix's AI-based recommendation engine, even if it consumes only a tiny part of Netflix's total Cloud usage).

³ Given the increasing importance of AI, there is a corresponding drive to increase the quantity of data. Varian (2018) suggests several methods for collection, including data scraping, finding public data repositories, entering data-sharing partnerships, or offering a service. The nature of data poses an interesting challenge to traditional strategy analyses, with their emphasis on resources that should be rare and inimitable (Barney, 1991), and whose apparatus is based on Ricardian notions of scarcity (Winter, 1995) and generally owned, when data only needs to be *accessed* instead.

⁴ To illustrate, the incidence of GitHub "stars" on TensorFlow (used to indicate GitHub members' appreciation) has grown at an annual growth rate of 63% since 2015, indicating GitHub's growing role as a collaboration hub.

One more historical particularity is worth considering here: the role of AI libraries and frameworks. These offer end-to-end ML tools that allow developers to build and train AI models; as such, they are a vital aspect of the division of labor in AI, and central to the diffusion of innovation in AI algorithms. AI libraries and frameworks attract talented developers to contribute to AI innovation, helping to address the problem of talent scarcity. They also entail economies of cost, time, and risk by “pooling” algorithms that have already been test-proven and/or peer-reviewed, and accelerate commercialization by integrating computing resources and industry data.⁵ AI libraries and frameworks were first established to explore cutting-edge research in AI (e.g., Torch); since then, they have more often been adopted and developed by tech giants (e.g., TensorFlow [Google], CNTK [Microsoft], and PyTorch [Facebook]).

For historical reasons, AI libraries and frameworks in the West have relied mostly on Open Source concepts: when “opening” an AI library, it’s crucial to attract users who will populate its shelves. Open Source was also central to the philosophy of many AI developers and researchers. But now that more established players have emerged, some tend to opt for “freemium”-like business models— notably in the U.S. For instance, TensorFlow is partially open-sourced by Google to attract talent to the platform, thereby remaining free to academics (TensorFlow Research Cloud). In return, academics are expected to use TensorFlow for programming, to share/publish research results, and help Google improve TensorFlow. On the other hand, enterprises must either meet their own costs for using TensorFlow, or pay to access TensorFlow Enterprise for improved versions of AI frameworks and consulting services. In China, AI libraries and frameworks are naturally developed for industry applications. Services such as Paddle-paddle by Baidu, Alibaba’s PAI, and Tencent’s TI all provide AI solutions for various sectors on a subscription model. As the industry structure is more fragmented, with more SMEs looking for off-the-shelf solutions, commercialization for incumbents is much stronger; this is the biggest difference between China and the U.S. in this regard. Also, AI libraries and frameworks in China emerged later than those in the U.S. (2014–15 vs. 2002), largely in response to demands arising from business use cases.

Another type of service has gained prominence recently. “No-code” AI platforms provide visual modules where core functionality is accessible through visual interfaces and pre-built integrations that can be use-case specific. This enables developers to build highly customized

⁵ Frameworks and libraries offer *building blocks* used by AI, which support higher-level functionality for cognitive tasks that are common to many applications, in particular for image processing and Natural Language Processing (NLP). They offer significant elements of pre-trained models, and enable modular approaches, more rapid development, and more robust implementation.

applications and systems at lower cost, without doing any programming in the conventional sense. Thus, more companies can leverage AI even if they can't recruit in-house developers.

2.3. AI in the *Cloud* and on the *Edge*

Finally, we come to the implementation of an AI application and supporting infrastructure for computing capabilities. Most data are stored and processed in bulk in the *Cloud*. Increasingly, however, businesses are implementing *Edge* devices—like IoT devices—that process data close to the source, complementing the *Cloud*. Edge computing minimizes latency and adds data pre-processing to devices, so more decisions or inferences can be made in “real time,” which reduces latency (the distance data must travel) and potentially the pressure on bandwidth (the amount of data that can travel in a packet). And by executing previously trained AI models from the *Cloud*, on-premise, *Edge* devices contribute to strengthening some security-sensitive applications such as facial recognition or autonomous driving—although the connection of these same devices to the outside world also raises a new security risk. *Edge* will become increasingly relevant as communication technology gets quicker and sensor technology improves.

3. AI's institutional structure of production and architecture

3.1. Understanding the key participants in the AI ecosystem

In the previous section, we described the technological setting in broad strokes. Yet, as the evolutionary approach (Nelson & Winter, 1982) suggests, technologies and competencies are rooted in specific organizations— as research in sectoral systems (Malerba, 2004) also confirms. Yet these organizations, and sectors overall, have boundaries that are set endogenously (Jacobides & Winter, 2005; 2012) as a response to competitive opportunities. Therefore, to understand the AI ecosystem, we need to focus on the “institutional structure of production” (Coase, 1937; Jacobides & Winter, 2005) that describes the division of labor between different participants. Per work on industry architecture (Jacobides, Knudsen & Augier, 2006; Pisano & Teece, 2007), we look at the evolving dynamics of AI producers and enablers. This gives us a lens for exploring the sectoral division of labor, the roles of different sector participants, and key rules. To do this, we need to understand how these different parties, more or less integrated into the AI ecosystem, monetize their advantage.

While a full description of the AI's division of labor would be a project in itself, it is important that we provide an overview of the different players, and how they engage with each other. Figure 3 below provides such a summary for the two parts of the overall

technology and IT stack (i.e., vertical segments): AI production (i.e., AI platforms) and AI enablement (i.e., supporting infrastructure).

Insert Figure 3 about here

This picture shows that the enablement and production stacks are mostly—though not entirely—controlled by Big Tech, which are integrated end-to-end, encompassing most stacks (especially Google, Amazon, Microsoft, Alibaba, and Tencent). These largely vertically integrated firms are motivated to encourage the production of AI (and trumpet its advantages), inasmuch as they stand to benefit from the improvement in what they offer: Amazon can improve its ability to target and sell (and possibly price) more advantageously; Google can increase its predictive accuracy, both within offerings like Gmail, Maps, and Search, and in the way they combine in a multi-product ecosystem; Facebook can improve its services, image recognition that facilitates complementary services and enhances customer lock-in, and also its ability to generate and sell advertising data; Microsoft can improve its software applications and business services. Hyperscalers also benefit from the upstream increase in the use of Cloud computing services they provide.

3.2. The AI enablement, production, and consumption actors, and their categories

While our focus is on the AI production and consumption landscape, we first need to address the AI enablement stacks required to produce and consume AI. Enablement is mainly composed of two layers (a.k.a. *bricks*, a term that we will use interchangeably, as does the industry). The first layer is hardware (sensors, chips, storage infrastructure, etc.) with significant competition in the chip/microprocessor environment (mostly led by Nvidia for AI applications, also including Intel, AMD, and Chinese players such as tech giant Huawei, and “unicorns” like UK-based Graphcore), as well as in sensors (e.g. Lidar for autonomous vehicles). The second layer is formed of data processing and management (e.g., companies or communities such as Cloudera, SQL, etc.), which, although not AI bricks themselves, can use AI to improve their products and are crucial to interfacing data inputs with AI production environments. Building high-quality data through data engineering and data labeling is becoming an important part of the process.

Focusing now on the AI production and consumption blocks, we see a number of different operating modes that appear to coexist. Overall, there are three ways of covering AI production needs (purchase, in-house production, or mixed supply), and two types of downstream uses (internal AI consumption or sale of AI solutions for clients to consume). In

general, companies fit neatly into the preceding categories, with the most overlap being in AI giants (e.g., Google, Alibaba etc.).

Insert Table 1 about here

AI giants (e.g., Google, Amazon, Alibaba, Tencent) have the capability to produce the AI they need for internal and external use. These companies operate every brick of the AI enablement and production sectors, and also consume what they produce in different parts of their business (e.g. Google search engine, which has a constant need for AI improvement, as its competitive advantage critically relies on its predictive power; ditto Amazon, with its ability to target customers and optimize logistics). Giants maintain market positions in every brick of AI by organically developing, partnering with, and acquiring leading companies (Google acquired AI start-up DeepMind, who successfully developed the AlphaGo program)—although future acquisitions might be affected by the recent shift in regulatory attitudes (Jacobides, Bruncko, & Langen, 2020). These companies, having benefitted from some advancements that relate to their own verticals, are also interested in finding other ways to monetize parts of the solutions they have produced, offering services to other players (e.g., infrastructure as a service, analysis as a service, data processing as a service, etc.).⁶

AI-powered operators/applications leverage AI in their day-to-day operations and their offerings. They use both AI giants' services and internal capacity to produce the AI necessary for critical functions/operations. Such operators include Facebook in social networking, Uber in mobility, Bio-N-Tech in healthcare, and PingAn in insurance. They have strong AI capabilities internally, use AI as a core aspect of their business model, and often enable it through AI solutions produced in-house. Given that AI solutions form part of their competitive advantage, these firms tend not to create revenues by selling production solutions developed in-house.⁷

AI creators have the capability to produce AI for external use (e.g., Accern, a no-code AI platform for financial services, and Nearmaps, a data analysis provider). These companies

⁶ It is important to stress that AI giants develop solutions and services that, geopolitics allowing, can be technically deployed across national borders and have a universal appeal. They sell their services to tech-enabled clients that may be either national or multi-local, or those who tend to require deep expertise within given verticals. Thus the emerging division of labor between telcos (which are national companies, even though they may belong to multinational groups), specialized service providers who work one market at a time (from delivery services to ride-hailing), even if they benefit from some global economies of scale), and AI giants (i.e., Big Tech) is reshaping the industrial landscape. AI giants are also acquiring some expertise in terms of application fields (e.g. Google's Verily venture in healthcare) but they do not aspire to cover end-to-end needs, as these require intense local engagement structure. Amazon, which draws on its e-commerce clout, may become an exception in this regard.

⁷ Unlike the AI giants, some of the AI-powered operators operate in, and focus on the particularities of, specific markets and geographies.

produce AI primarily for specific third-party use-cases, and less for their own use. They largely rely on tech giants' platforms and services to obtain generic AI solutions, which they subsequently improve and/or customize to their clients' needs.

AI traders/integrators buy and sell AI solutions or use cases, adding commercial and marketing services (e.g., bundling, repackaging, branding) or supporting integration with the client's ecosystem, but without making any AI improvements (e.g., translation companies using Google Translate; CRM consultants integrating Salesforce or MS offers, etc.).

AI takers require AI solutions (off-the-shelf or customized) to enable critical business functions (e.g., digital natives focusing on a narrow AI value-add element and outsourcing most of the AI production they require, or traditional companies with limited internal AI capabilities). They are often incumbent or traditional companies looking to transform by using AI solutions, or startups without the ability or funds to develop in-house. These companies can live with a standardized AI solution or pay more to have it customized, but they can't develop it internally. Interestingly, although they do not produce their chosen AI solutions internally, they will usually improve endogenously over time, because of the learning loop we discussed in section 2. To source AI, one solution is to enter into a partnership with AI giants or AI-powered applications, giving takers access to additional data in exchange for cheaper technology. Another solution is to buy AI at "full price," from either subsidiaries of the Big Tech firms or vertical specialists, in some of the categories above.

3.3. The economics and sectoral dynamics of AI

Having considered the different firms in the AI production and enablement stacks, we should also consider the underlying economics. The evolutionary dynamics of this complex ecosystem are driven by economies of both scale and scope. The former relates to the cumulative advantage of the tech giants, and imposes significant barriers to entry. The latter relates to their abilities to grow laterally, entering new verticals. The availability of large amounts of data (the "core input" of this "fifth wave of development," per Freeman and Louca's (2001) chronology) is central to both.

First, the enablement and production stacks are characterized by massive capital intensity, and potentially economies of scale, possibly enhanced by economies of learning. These are areas where scale begets learning through the accumulation of data and increases competitive advantage to such a degree that a few firms—such as Google, AWS, and Microsoft—have emerged under the term "hyperscalers." Initially, they scaled to serve their own needs, but

increasingly compete by making computing commercially available *on demand*.⁸ In terms of the AI production side, there are a number of scale-intensive areas. These algorithmic models (that can be found on PyTorch hub or TensorFlow hub), especially in the growth areas of ML (such as natural language processing), seem to lead to very significant economies of scale. As Sharir, Peleg, and Shoham (2020) note, “Just how much does it cost to train a model? Two correct answers are ‘depends’ and ‘a lot’ ... Exact figures are proprietary information of the specific companies, but one can make educated guesses ... the total ... price tag [of one specific model that was tested] may have been \$10 million.” Given that a number of these models need to be produced for any one AI predictive model, this stack favors larger players. This means that we might soon have a setup whereby a few firms (e.g. those with libraries and platforms) do all of this work, and allow an ecosystem of co-specialized complementors (per Jacobides et al, 2018) to support them by fitting models to applications. As Sharir et al (2020) observe, “Not many companies – certainly not many startups – can afford this cost. Some argue that this is not a severe issue; let the Googles of the world pre-train and publish the large language models, and let the rest of the world fine-tune them (a much cheaper endeavour) to specific tasks. Others (e.g., Etchemendy and Li [2020]) are not as sanguine.”

One final issue arises in relation to economies of scope. We know that access to *data* is critical for AI, and that firms who have rich data will be incentivized to invest in and leverage AI. We also know that larger databases reduce the computational cost of “training” models—or, equivalently, that they increase predictive accuracy as datasets grow (see Kaplan et al, 2020), meaning that actors with bigger datasets will have better returns or lower costs in developing AI models. The use of AI reinforces the importance of Cloud as an industry. Many Cloud providers also own AI platforms, allowing them to control a large portion of the industry.

3.4. AI sectoral dynamics driven by the Edge, and upstart power

The growth of the Edge in terms of relative importance has led to the creation of an interesting web of activity, which has attracted the entry of both *de novo* and *de alio* players (Sosa, 2013) from a variety of backgrounds, such as real estate, hardware, connectivity services, and industrial goods. While major tech players such as Siemens and Bosch are

⁸ Such Cloud services have, for instance, been Amazon’s largest contributor to growth and profits over the last decade. To give a sense of the scale at the plant level, investments in single centers by hyperscalers currently range between \$1–3 billion (see Miller and Laird, 2019), and firms like Microsoft, further anticipating benefits from multiple such centres, have changed the core of their business model to enable the capital expenditure necessary for being a key hyperscale provider. While we do not have details of either market power or margin, some security concerns have recently been raised (see Zuo & Goines, 2020).

leveraging their core strengths to build industry-relevant Edge IoT products⁹, start-ups mostly focus on application and analytics software on-device, and predicate their Edge technology on specific use cases. Figure 4 below illustrates the key players in this domain.

Insert Figure 4 about here

Edge and IoT devices are rarely created by the same companies that own Cloud computing capabilities (one exception being smart home devices such as Google Home). This raises technical challenges in terms of integrating the two. In December 2019, Google, Amazon, Apple, and the Zigbee Alliance formed Project Connected Home to create a standard for smart home device compatibility.¹⁰ This project group seeks to simplify manufacturing and increase options for consumers, and thus enable modular co-development of Edge and IoT devices in this area (Baldwin & Clark, 2000). This sort of industry standardization will likely reinforce the position of Cloud service providers, at the expense of Edge and device providers—typifying the types of strategic challenges that need to be modeled in the world of AI (see Baldwin & Woodard, 2009 and Jacobides & Tae, 2015 for broadly similar analyses).

Possibly as a reaction to such increasingly concentrated structures, especially in the Cloud (but also potentially the Edge), the AI community is increasingly engaging in and supporting platforms that share development costs without being wedded to one of the Big Tech ecosystems, such as huggingface.co and rasa.com. Also, while the dominance of hyperscalers is absolute, there appears to be clear space for ventures that work alongside them, maintaining the levels of “speciation” (Saviotti, 2005) in this otherwise highly concentrated ecosystem. In 2019, the AI venture market surpassed \$27 billion in total volume, with over 2,235 deals (CB insights), having grown at an annual rate of 29% since 2015.¹¹ The total number of deals exceeded the 9% annual growth of deals by hyperscalers including Google, Microsoft, and Amazon. In addition, the total monetary value of AI investments made by hyperscalers (Google, Microsoft, and Amazon) accounted for around 14% of total market investments till 2018.¹² In rare instances, new entrants managed to outcompete hyperscalers in particular segments: Snowflake, a data management service provider since 2012, has grown to be a market leader in a key product area dominated by Amazon.¹³ Although hyperscalers sustain

⁹ Bosch’s Home Connect sensor range includes an array of IoT-enabled devices and sensors for smart home and industrial uses, such as smart washing machines, ovens, thermostats, cameras, etc.

¹⁰ https://zigbeealliance.org/news_and_articles/connectedhomeip/

¹¹ <https://venturebeat.com/2020/01/22/cb-insights-ai-startup-funding-hit-new-high-of-26-6-billion-in-2019/>

¹² See <https://www.oecd.org/going-digital/ai/private-equity-investment-in-artificial-intelligence.pdf> ; CB insights: <https://www.cbinsights.com/research/facebook-apple-microsoft-google-amazon-ai-investments/> ; <https://www.techrepublic.com/article/the-10-tech-companies-that-have-invested-the-most-money-in-ai/>

¹³ <https://finance.yahoo.com/news/snowflake-gives-investors-rare-opportunity-100000455.html>

outstanding competitive advantage due to access to data, top talents, and computing resources, the AI market is large and growing fast, with competition furthering innovation. So, even with further consolidation at the top possible, the entry of newcomers and the vitality of the ecosystem seem assured.

4. AI in Action: From Production to Application

Having looked at all the facets of AI production, and the firms involved in it, let us now move to AI use. AI, and ML in particular, has made great strides because of technological advances in some key areas, which have greatly facilitated downstream applications. In vision recognition, for instance, developments in the last decade have been likened to the Cambrian explosion 500 million years ago, when trilobites and other sea creatures developed vision, leading to a proliferation of life forms (Pratt, 2015). These technologies have enabled a massive improvement, and significant uptake, in factory automation and AI medical diagnoses. They have also enabled a number of services to be offered via social media platforms like Facebook, which have both increased customer engagement and allowed for complementors to leverage their data and find new ways of monetizing their advantage.

4.1. Upstream AI provision and downstream Big Tech demand

Looking at usage patterns, we see that today's Big Tech firms have played a key role in funding and promoting the development of AI, which is consistent with their (downstream) business models. This observation is quite clearly at odds with the concerns about AI underinvestment raised by those who see AI as a GPT. Moreover, for some (the hyperscalers), AI growth also leads to massive uptake of their upstream services in Cloud computing. This explains, in our view, why AI does not suffer from under-investment. Given the use of AI by Big Tech, much of the required investment has been directly funded by them. Indeed, corporate departments are publishing more papers than scholars—an extraordinary situation compared to any other field of science. Entire new sub-fields, such as Federated Learning, have essentially been created by GoogleAI. This is a crucial observation, particularly in the context of growing concerns about the declining role for basic research in corporate R&D (e.g. Arora et al, 2018). Here, we have a different challenge, with research dominated by a handful of firms. The role of a strong science system is central to any evolutionary account (e.g. Metcalfe, 1994). Science and technology research is a key engine to generate novelty within systems that tend to follow cumulative dynamics (such as those enabled by data accumulation, as discussed above). Perhaps disconcertingly, the role of Big Tech firms (especially Big Tech firms in the U.S.) in driving the AI research trajectory is growing stronger and stronger. Consider, for instance, the number of papers in NeurIPS (Neural

Information Processing Systems) and ICML (International Conference on Machine Learning), the two premier conferences on ML hosted in 2019. Google had the most papers by far, followed by Stanford, MIT, CMU, and UC Berkeley, then Microsoft at number six and Facebook also appearing in the top 10¹⁴.

Big Tech's focus on AI has been so strong that there is increasing concern about the viability of publicly funded research in particular fields of computer science. Until 2004, an inflection year for AI, no AI professor had left academe for machine learning; between 2014 and 2018 Google/DeepMind hired 23 tenured/tenure track faculty, Amazon hired 17, Microsoft hired 12, and Facebook, NVIDIA, and Uber hired seven apiece. This is probably only the tip of the iceberg, and similar moves are anecdotally known to happen in China. With AI being increasingly seen as vital to both corporate success and economic growth, the explosion of activity has also led to a funding "arms race" between national governments (despite the fact that the prime beneficiaries appear to be a very specific type of firms and their ecosystems). The EU has pledged €24 billion, while China's target is \$150 billion by 2030. Generally, these trends highlight the centrality of the discussion about skills and capabilities behind the emergence of the AI ecosystem. What capabilities the tech giants are developing we can only infer from the products and services they launch. Secrecy, rather than patenting, remains the preferred strategy to protect their research findings.

4.2. Attributes of (successful) AI adopters

The firms that have deployed AI have some very unusual characteristics, as Iansiti & Lakhani (2020) argue. They have a different operating model; they are driven by data; they have redefined processes to put AI at the core; they engage in experimentation and make decisions in real time; they perform extra granular forecasting; and they learn proactively from the reaction of their customers, employing real-time experiments in their offerings and evaluating the data. That said, hyperscalers with advantages in almost all the aforementioned dimensions bear almost zero marginal cost on deploying AI to larger scale.¹⁵ Thus, their effective use of

¹⁴ Calculated by number of publications on NeurIPS and ICML. For each publication, each participating organization will be scored as number of authors from that organization divided by total number of authors for the publication. See <https://chuvpilo.medium.com/ai-research-rankings-2019-insights-from-neurips-and-icml-leading-ai-conferences-ee6953152c1a>

¹⁵ As Iansiti & Lakhani (2020) note in the summary of their *HBR* article, summarizing their book, "Rather than relying on processes run by employees, the value we get is delivered by algorithms. Software is at the core of the enterprise, and humans are moved off to the side. This model frees firms from traditional operating constraints and enables them to compete in unprecedented ways. AI-driven processes can be scaled up very rapidly, allow for greater scope because they can be connected to many kinds of businesses, and offer very powerful opportunities for learning and improvement. And while the value of scale eventually begins to level off in traditional models, in AI-based ones, it never stops climbing. All of that allows AI-driven firms to quickly overtake traditional ones. As AI models blur the

AI is predicated on several organizational practices that are prerequisites for AI to have an impact. These have been noted by Brynjolfsson et al (2019), who, drawing on earlier work on IT adoption more broadly (Brynjolfsson & Hitt, 2012), hypothesize that the lack of these complementary investments is what impairs the impact of AI—at least in terms of productivity statistics.

We broadly agree with this thesis, and would qualify it. The micro-level and behavioral evidence clearly suggests that adopting AI in isolation may be fashionable and seen as progressive, but it is a real challenge for companies to see a return on investment. This message comes through very clearly from all consulting reports on AI. For instance, BCG Henderson Institute and MIT conducted a study (Ransbotham et al, 2020) showing that although more than 50% of companies are deploying AI, only 11% report significant benefits from implementing it. These findings suggest companies have far to go in order to harness the benefits of AI. For instance, even within companies who invested in building foundational capabilities—AI infrastructure, talent, and strategy—only 21% achieve significant financial benefits. However, when firms focus on what BCG calls “organizational learning with AI”—i.e. implementing AI at scale while explicitly focusing on human/machine collaboration—the likelihood of realizing significant financial benefits leaps to 73%. This illustrates the challenge facing companies, given the inherent complexity of the technology and the effort and time required to redesign the organization around AI (Ransbotham et al, 2020; Tambe et al, 2019).

Clearly, only the firms that are already proficient in managing operations, using data, and engaging with customers will be able to generate value from AI deployment (Bock & von Wangenheim, 2019). Our qualification to the thesis of Tambe et al (2019) is that we do not think that it is merely a matter of time for these changes to “trickle through.” We believe that some of the productivity adjustment may happen through inefficient or “non-digital-friendly” firms eventually losing market share or being out-competed. In other words, we anticipate significant and systematic inequalities when it comes to the impact of AI—in terms of both who *uses it* and who *benefits*.

This observation is confirmed by the sole systematic and rigorous academic study of AI adoption at the firm level we know of, which draws on data on job postings in AI from Burning Glass Technologies to proxy the deployment of AI. This paper finds that AI

lines between industries, strategies are relying less on specialized expertise and differentiation based on cost, quality, and branding, and more on business network position, unique data, and the deployment of sophisticated analytics.”

investments are concentrated in the top tercile of firms in each sector (measured by performance), and furthermore that the most profitable and effective firms are those who benefit the most (Babina, Fedik, He, & Hodson, 2020).¹⁶ These findings are consistent with what an evolutionary account would expect. Investments in technology per se do not drive performance; they must be complemented by investments in managerial and organizational capabilities that support the continuous transformation of ideas into products and services (e.g. Nelson & Winter, 1982; Cefis & Ciccarelli, 2005), and better firms tend to be endowed with such superior (dynamic) skills (Teece, Pisano, & Shuen, 1997; Bloom & Van Reenen, 2010).¹⁷ This complementarity reinforces the cumulative, path-dependent nature of the evolutionary processes we observe in the AI ecosystem.

4.3. AI adoption, data access, business models, and complementor communities

The use of AI is predicated not only on performance, but also on access to data, as Clough and Wu (2021) have recently pointed out. This is another crucial distinction between Big Tech firms and the rest. Firms such as Google, Facebook, and even Apple distinguish themselves by having business models that rely on an extraordinary rich set of information on their customers (see Jacobides, Bruncko, & Langen, 2020, for a detailed analysis). The same applies to Amazon, and even more to Chinese Big Tech players such as Alibaba and Tencent. Whether these firms own the data or not (see Varian, 2019), they surely have the right to use it, which makes AI an important tool. This is not necessarily true of other firms, which raises a more general point: analyses of AI and its use might be conflating the *technology* of drawing lessons or predictions from data with the *opportunity to use such data*. Firms with data access are also firms with AI investments; and their productivity and, particularly, profitability differences may be due to both.

The data used for AI is not necessarily owned, but merely accessed (Clough & Wu, 2020). Big Tech firms, for instance, ensure that their ecosystems are structured in a way that allows them to benefit from their own data, but also that of their complementors; Google and Facebook access real-time information on user behavior from software that connects to their

¹⁶ As they note, “larger firms, in terms of both sales and market share, are more likely to invest in AI, consistent with the evidence by Alekseeva et al. (2020). Furthermore, AI investments are stronger among firms with larger cash holdings, higher mark-ups, and higher R&D intensity... Firms that invest in AI grow more. Specifically, a one-standard-deviation increase in the share of AI workers based on the resume data corresponds to a 15.6% increase in sales, a 15.2% increase in employment, and a 1.4 percentage point increase in market share within the 5-digit NAICS industry... *the positive effects of AI on firm sales growth are concentrated in the most ex ante productive firms, with large positive effects for firms in the highest productivity tercile in 2010 and small and insignificant effects for firms in lower terciles.*” (emphasis added)

¹⁷ Consistent with Iansiti and Lakhani (2020), we regard performance as an incomplete proxy for how effective and digitized a firms’ processes are.

own with minimal to no coding integration (e.g., APIs) as the firms optimize for between-device compatibility and intra-device communication protocols (see Jacobides, Bruncko, & Langen, 2020). On the flip side, Big Tech firms have also shared data that allows for AI models to be trained (e.g., the Open Images Dataset), as Hal Varian (2020), the Chief Economist of Google, notes in his review of AI.¹⁸

A related observation is that while Big Tech firms' business models benefit directly from AI and create value for the entire sector (e.g., social media or digital marketing), they also generate business for their complementors. As such, the growth of digital ecosystems (Adner, 2017; Jacobides, Cennamo, & Gawer, 2018) goes hand-in-hand with the growth in AI. Firms that collaborate with Big Tech find ways to benefit from Big Tech's data. This is consistent with the findings of Babina et al (2020), who find that *sectors* that use AI benefit overall, and suggests that Big Tech operate as “kingpins” (Jacobides & Tae, 2015)—firms that create benefits for themselves and their segment (or, here, their complementors) by advancing technology, while skewing the distribution of profits. This leads to a feedback loop between technological edge, resource accumulation, and market dominance.¹⁹

Looking ahead, and drawing on our analysis, it looks likely that the future of AI innovation and leadership might require much more industry specialization than today, which could shift power from Big Tech to vertical-specific ecosystems that encompass both Big Tech firms and traditional companies. (The Online Appendix offers some evidence on the current patterns of downstream AI use, and expectations about the future.) As technology accelerates, it will open up new avenues for innovation and data collection (e.g., better sensor technology, faster processing, and no-loss storage), enabling further AI applications and innovation. This acceleration will continue to fragment the AI landscape and create new and more specialist use cases requiring increasingly refined techniques (e.g., autonomous surgeons). For general AI use cases, there is still a need for ML scientists to “manage the tail” of data, as algorithms' capabilities are still limited in Edge case data (e.g., an autonomous vehicle capturing a video of a human gesturing on the side of the road, which is understandable in human context).²⁰

¹⁸ Data sources, in addition to using the by-product of operations, include offering a service, web scraping, relying on Cloud providers (who host and provide data), government sources, pooled / purchased / compiled data, etc.

¹⁹ As Babina et al (2020) say, “We find that industries that invest more in AI experience an overall increase in sales and employment within the sample of Compustat firms ... AI investments not only spur industry growth, but also increase industry concentration. A one-standard-deviation increase in the share of AI workers based on resume data increases sales by 17.3% in the top tercile of initial firm size, 4.3% in the middle tercile, and 0.0% in the bottom tercile.”

²⁰ <https://www.zdnet.com/article/weird-new-things-are-happening-in-software-says-stanford-ai-professor-chris-re/>

5. Learning from international differences: The AI Systems of China, the U.S., and the EU

As our discussion above has alluded to, while AI may have some common attributes across different sectors, there are also some significant differences. The evolutionary aspect of AI ecosystems is partially shaped by their environments, which vary widely across geographic areas in terms of commercial, academic, regulatory, political and cultural background. And this matters a great deal, as past work on National Systems of Innovation has made abundantly clear. The very nature of the “triple helix” (Etzkowitz & Leydesdorff, 2000)—i.e. the way that government, institutions, and firms interact—affects the way things are organized. Given that the key areas for AI development are China, the U.S., and the EU, we will focus on them here, and in Online Appendix 3.

5.1. A tale of diverging triple helixes

First, let us consider what appear to be the “static” differences. These are summarized in Table 2 below, which originates from a 2020 BCG survey of large companies. The international contrast in responses is an indicator of differences in attitudes: according to respondents, 86% of end users of AI in China generally trust the AI solution’s decisions, while only 45% of European users and 39% of American users do so. Accordingly, if we take executives’ views as proxies for their compatriots, AI users in China have greater faith in AI, and are more patient with it. Over 80% of executives surveyed think that end users in China believe that AI improves business outcomes and understand the inner workings and limitations of AI systems (35%/54% for U.S.; 28/62% for EU). The level of public understanding and interest in AI is another factor influencing the business perspective. Respectively, 87–89% of companies in China are interested in AI because customers and suppliers are asking for AI-driven offerings, while only 61–72% of U.S. companies and 63–69% of EU firms share their reasoning. The question then is: how come China is so positive about AI? This is where some of the triple-helix dynamics, as well as the co-evolution of the sector and particular firms comes in.

Insert Table 2 about here

Unlike the U.S. and the EU, China seems to have benefitted from its own “Sputnik moment” (Lee, 2018), when AlphaGo defeated Korean player Lee Sedol at the traditional board game of Go in 2016. Former Google China president Kai-Fu Lee (2108) explains that the five-game series “drew more than 280 million Chinese viewers and lit a fire under the Chinese technology community.” In May the following year, the defeat of Go champion Ke Jie even accelerated Chinese actions on AI. Less than two months after the game, the Chinese

government issued the Next Generation Artificial Intelligence Development Plan²¹, which called for greater funding, research and innovation, and national cooperation for AI, and outlined ambitious goals to reach the top tier of AI economies by 2020, achieve major new breakthroughs by 2025, and become the global leader in AI by 2030, with very significant funds committed centrally. These government initiatives were matched by changes in businesses and academe, leading to significant dynamism in AI.²²

As Lee (2018) notes, by 2017, Chinese venture capital investment in AI made up 48% of AI global funding, surpassing the U.S. for the first time. Chinese Ministry of Education figures reveal that undergraduate majors (degrees) related to AI,²³ which totaled just 64 in 2016, jumped by 328 in 2017 and a further 510 in 2018.²⁴ As Arenal et al (2020) note, Chinese central government support through its strategic plan and support of AI clusters (including universities and enterprises) quickly paid off, while in the U.S. there was less central control, and firms were left to their own devices. Indeed, Big Tech firms, which were mostly based in the U.S., took the initiative on AI investment, and a number of ventures duly emerged, albeit without much planning (as we explained in Sections 3 and 4). The U.S. government did not consider that AI (or AI infrastructure) needed to garner such support—leading American academics to call for greater government involvement, without such a reliance on hyperscalers.²⁵

These differences also manifest themselves in terms of industry architecture: the rules and roles for the division of labor, but also how key firms form their ecosystems, and how they engage with their complementors. With a better understanding of AI (development cost and

²¹ See details on the Chinese 2017 plan in <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/translation-chinese-government-outlines-ai-ambitions-through-2020>

²² The target for 2020 was that the overall level of technology and application of AI of China should have catch up with leading countries in the world. The market size of AI in China should have reached \$21 billion by then. For future targets, by 2025, China should make great breakthroughs in basic theories of AI, and some technologies and applications achieve world-leading level. The market size of AI in China should reach \$57 billion. And by 2030, China's overall AI theories, technologies, and applications achieve world-leading levels, and China should become the major AI innovation center of the world. The market size of AI in China should reach \$143 billion. Source: New Generation Artificial Intelligence, The State Council of the People's Republic of China. See http://www.gov.cn/zhengce/content/2017-07/20/content_5211996.htm

²³ AI-related majors include artificial intelligence, smart science and technology, robot engineering, smart manufacturing, data science and big data, and big data management and application.

²⁴ Source: Ministry of Education of the People's Republic of China. See http://www.moe.gov.cn/srcsite/A08/moe_1034/s4930/201703/t20170317_299960.html

²⁵ Etchemendy and Li (2020), for instance, make an impassioned plea to reconsider the current status quo, in which, as they note, “[P]ublic researchers’ lack of access to computer power and the scarcity of meaningful datasets, the two prerequisites for advanced AI research ... threatens America’s position on the global stage,” and argue for a National Cloud Service in the USA, introducing an additional dimension not only of public vs. private infrastructure, but also of geopolitical clashes that seem to be shaping the views of what are the appropriate industry architectures. In a time of a growing U.S.–China (and, potentially, EU) technological conflict, this raises yet another dimension of policy concern.

time of; resultant benefits), Chinese companies are more dedicated to AI adoption with strong leadership support and cross-functional teams created up front to support to the development of AI solutions.

Political context also contributes to the vibrancy of AI ecosystems. For instance, in China, tech giants are encouraged by governments to establish AI libraries/platforms to enhance ecosystem partnership and allow SMEs to access AI technology at a lower cost. Thus, per the government's request in 2017, Tencent was chosen to lead AI innovation in computer vision for medical imaging, Baidu for autonomous driving, Alibaba for smart cities, SenseTime for facial recognition, and iFlytek for voice intelligence. The Chinese authorities further expanded the "national task force" into 15 companies in 2019, asking them to export their tech capabilities for industry incumbents through collaboration in open data, algorithms, and models, theoretical research, and applications (especially for SMEs and start-ups)²⁶. These tighter links between government and business mean that Chinese tech companies orchestrate ecosystems with incumbents within various industrial sectors to support those sectors' transformation with AI. This offers tech companies additional data and access to new markets in return for their provision of digital and AI capabilities, products, and services, and creates distinct industry architectures.

In contrast, while U.S. tech giants also form some partnerships with incumbents from time to time (for instance, Google-Waymo with Fiat Chrysler for AVs), this is much less common than in China. Indeed, the U.S. AI transformation is mostly driven by numerous waves of incumbents being replaced by fast-growing tech companies that are vertically specialized and fuelled by abundant VC funding. In Europe, on the other hand, where the VC industry is less mature, for all the emphasis on regulating AI and setting moral and ethical boundaries, state (or EU) business interventionism is more limited. This leaves the initiative to individual companies like Siemens, who launched industrial challenges to recognize leading AI companies such as Top Data Science and embed their solutions into Siemens' own ecosystem (Siemens IoT platform, MindSphere),²⁷ or the recent "Software République" initiative led by Renault²⁸ to join forces with four large French companies to create a new ecosystem for intelligent and sustainable mobility. This places considerable onus on firms with little experience of creating broad alliances or building their own ecosystems. Figure 5 below

²⁶ Source: Ministry of Science and Technology of the People's Republic of China .See http://www.gov.cn/xinwen/2019-08/04/content_5418542.htm

²⁷ See <https://topdatascience.com/2019/06/12/siemens-and-tds-announce-collaboration-in-ai-for-bio-based-industrial-processes/> and <https://hitinfrastructure.com/news/siemens-ui-healthcare-to-enhance-medical-imaging-technology>

²⁸ See <https://en.media.groupe.renault.com/news/groupe-renault-atos-dassault-systemes-stmicroelectronics-and-thales-join-forces-to-create-the-software-republique-a-new-open-ecosystem-for-intelligent-and-sustainable-mobility-a31a-989c5.html>

provides a visual summary of the key international differences, including regulatory attitudes and the relative roles of firms, governments, and academe. More details are provided in Online Appendix 3.

Insert Figure 5 about here

These differences between contexts are also manifested on the diverging practices that link firms and developers. As we noted in Section 3, Chinese firms develop libraries for other entities to use—but unlike U.S. Big Tech and hyperscalers, they do not give them away for free. This is partly due to a different trajectory – with the “open source” movement being much more prevalent in the U.S., and by extension in Europe—and also because firms find different ways to generate benefits for themselves.

5.2. From National AI Systems to Global Firm-Based AI Ecosystems

These national differences, important as they might be, do not imply that all dynamics happen within countries. Countries (or, in the case of the EU, regional groups with significant resources and authority) set the rules and shape the ecosystem locally—which can be seen, for instance, in the European Commission regulations on data sharing. However, some of these activities are global—precisely because of the extreme economies of scale and reuse, and the ability to learn from more data. So, for at least *some* parts of the AI sector, we have both local and global dynamics, inasmuch as some of the players have a strong interest in leveraging their advantage on a global scale. The ease with which firms from different geographies can do this differs, with U.S. firms having moved much more aggressively in terms of their global scale than Chinese firms, and EU firms having mostly kept it small. However, there is an increasing part of the AI ecosystem that is becoming global, both in terms of the “on-demand hyperscalers” and their attendant AI services (which are facilitated by judicious location choices around the globe), and because the expertise and research can be leveraged more broadly.

Some firm-specific ecosystems span the globe, providing an interesting new dynamic whereby, in addition to National Innovation and Sectoral systems, we have global AI ecosystems that cut across traditional divides. To quantitatively illustrate the connection between orchestrator and ecosystem members in AI, Rock (2019) finds that following the release of Google’s TensorFlow, the value of firms in the AI sector jumped. To illustrate Google’s global ambitions, and focus outside its home market, consider the geographical breakdown of programmers using Tensorflow.²⁹ In addition to the c. 380,000 contributors

²⁹ Data on Members of the TensorFlow certificate network were gathered on April 6, 2021.

worldwide, there are 1,195 Premium Contributors, of which only 370 are in North America, 1,168 in Asia (excluding China, which bans Google), 347 in Europe, 30 in South/Central America, and 20 in Africa.

In terms of what we expect, it is hard to predict, as in addition to technological uncertainties, geopolitical uncertainties also play a role. As noted above, the democratization of AI is enhanced by the rapid proliferation of AI services and libraries offered by Big Tech firms. Hyperscalers also stand to benefit significantly from the growth in Cloud computing services. This, in turn, generates significant demand and inequality in spending, as well as aggravating global warming—a challenge that is becoming increasingly clear (Dhar, 2020).³⁰

Downstream demand for AI will continue to be encouraged, though, given the current incentives. Big Data firms like Amazon and Facebook, but also Google and Microsoft, are offering some basic AI functionality in their core products, from the recommendation engine on Amazon marketplace to priority e-mail suggestions or end-of-sentence propositions in Gmail. Growing excitement at state level is fueled by an expectation of productivity gains, economic growth, and a fear of losing out in a geopolitical fight for technological supremacy. Some players are also offering significant support to firms that are considering the use of AI. In China, in particular, Big Tech firms that are both hyperscalers and also have large ecosystems of their own are proactively supporting firms not only to digitize, but also to employ AI. This may have the added benefit of ensuring that Chinese Tech ecosystems aim create a more equitable set of complementors, thus cementing their positions as kingpins, who can capture a greater part of the total value-add (Jacobides & Teng, 2015).

6. Leveraging the AI dynamics to apply and extend the “roots and branches” of evolutionary dynamics

Beyond the phenomenological interest in AI, and the use of evolutionary tools to comprehend its nature and dynamics, what can we take away from this paper, methodologically speaking? This section takes a step back to consider how the sectoral dynamics we analyzed not only apply, but also contextualize and extend evolutionary tools, and what they show us about the “roots and branches” of evolutionary theory, and how they relate.

6.1. How the AI sector case study can inform existing evolutionary tools

Our analysis of AI brings up some interesting observations inspired by research on “technological regimes” (Malerba & Orsenigo, 1999; Breschi et al, 2000)—that is, the

³⁰ To illustrate, training a single big language model generates around 300 tons of carbon dioxide emissions—as much as 125 round-trip flights between New York and Beijing.

technological conditions that determine whether small or large firms drive the creative process in a sector. The tension is between two settings. In the first, small entrepreneurial firms come up with new ideas before growing and becoming dominant, only to be deposed by a subsequent wave of “creative destruction” in a process dubbed “Schumpeter Mark I” (in reference to Schumpeter’s (1912) early work). This is contrasted with “Schumpeter Mark II,” where large firms have internalized the innovation process, in what Schumpeter (1942) described in his later book. In terms of these regimes, AI offers both an application and the opportunity to qualify the framework.

Let us first apply the technological regime framework to our setting. Which regime emerges is related to four key drivers: technological opportunities, the appropriability of innovations, the cumulativeness of technical advances, and the properties of the knowledge base. The core issue, in our view, is appropriability. To give a concrete example of questions that an evolutionary framework allows us to approach: was Google technologically unavoidable? If Google had never existed, would another “Google” have emerged to fill the technological and strategic void? Interestingly, Google’s own founders reveal that they had different options to choose from. In their well-known 1998 article, they wrote: “[W]e believe the issue of advertising causes enough mixed incentives that it is crucial to have a competitive search engine that is transparent and in the academic realm” (Brin & Page, 1998). We all know that things went a very different way. Yet, it is important to note that Brin and Page themselves framed the problem in terms of incentives, not technological requirements. It was an economic and managerial choice that led to this strategy—which, in turn, shaped the technological trajectory.

This choice was made possible by the fact that data could indeed be used and potentially appropriated, whether due to cultural and political acceptability (e.g. China), or due to the lack of clear regulations (e.g. in the U.S., or the EU before the GDPR). This was a crucial moment in the evolution of the AI sector, because the appropriation of data led to powerful network externalities that, in turn, made Big Tech firms possible. Once network externalities kicked in, vertical integration ensued.³¹ In parallel, given the modular features of parts of the digital infrastructure, specialization in downstream applications emerged as a viable technical and business solution. Hence, we observe not the dominance of Schumpeter Mark II-type firms alone, but rather their coexistence and interaction with smaller, specialized Mark I firms.

³¹ This is a hypothesis that a new generation of history friendly models (Malerba et al. 2016), focused on AI, could explore.

In addition to technological regimes being a useful tool to help us understand the AI sector, the regimes themselves may be updated as a result of this detailed analysis. We find in particular that AI, like other digital sectors, exhibits an unusual and distinct pattern, where success and innovation is not the province of *either* large or small firms, but rather comes as a result of collaboration between both forms. This variant of “Mark II,” which we might call “Mark IIa,” relies on large firms (here, Big Tech and other digital giants) who enlist a host of complementors and ecosystem partners and regenerate themselves while strategically managing their partners (e.g., through library provision and terms of engagement). More important, this is a model where the renewal and success of large firms happens through acquisitions—a pattern we see often in this sector. Our evidence is consistent with a positive feedback loop where tech giants poach the best talent from academe, as well as funding, which facilitates both their differentiation (or “competition on merit,” as antitrust scholars call it) and ecosystem lock-in and quasi-insurmountable obstacles to competition, enhancing their profits, rents, and role in the economy (see Jacobides & Lianos, 2021). This affords them the financial resources to engage in acquisitions that can nullify potential competitors.³²

6.2. What the Study of AI Shows Us about the Evolution of Evolutionary Thinking

This detailed case study of the AI universe can showcase the value-add of an evolutionary account—but it also illustrates how the roots and branches of evolutionary analysis combine to shed light on these fascinating dynamics. This is the idea behind Table 3, which shows both the evolution of this thinking, and how it relates to AI.

The evolutionary approach has consistently acknowledged the importance of adopting a systemic view of change and innovation, at different levels of analysis. For example, National Innovation Systems (NIS) literature has always recognized the central role of governments in shaping the evolutionary patterns of technology-based competition (e.g. Freeman, 1995; Lundvall, 2007; Nelson, 2020). The sectoral systems approach provides a natural foundation for our work, with its emphasis on the interplay of heterogeneous actors, capabilities, and institutions as the engine for innovation and change (e.g. Malerba, 2004). Work on large, technical systems has traditionally urged scholars to grasp the inner complexities of technologies in detail, in order to identify where and how strategy and agency have room to intervene and steer the evolution of the system (e.g. Hughes, 1983; Brusoni et al., 2001).

³² It would be useful to have a systematic analysis of the patterns of acquisitions in AI, and the absorption of top faculty talent, bundled with the growth of entrepreneurial firms usually seen as complementors and not competitors to the dominant players. Such an analysis would better describe the entrepreneurial regime and anticipate the evolution of the sector and its competition dynamics.

Beyond these “roots,” a number of intellectual progeny of the evolutionary view have emerged to expand, augment, and qualify these seminal contributions in the last few years, including work on triple helix (Etzkowitz & Leydesdorff, 1995); industry architectures (Jacobides & Winter, 2005; Jacobides et al, 2006); business and innovation ecosystems (Adner, 2017; Jacobides et al, 2018); and digital platforms (Gawer & Cusumano, 2014; Parker et al, 2016; Boudreau, 2010). Each of these approaches, which we see as the “branches” of evolutionary theory, helps shed light on a particular aspect of the empirical reality, and collectively provide a more robust framework—which itself evolves.

Insert Table 3 about here

Table 3 captures the focus of the various evolutionary approaches we have built on in our analysis (see bottom row), and endeavors to trace which recent research streams they have generated (left to right), along with the slices of empirical reality that each one can capture. For example, recent work on business and innovation ecosystems builds on ideas from the sectoral system of innovation approach, adding a focus on the role of organizations that aim to develop new functionalities within an established technical architecture (i.e. complementors). Work on large, technical systems is related to current discussions about digital platforms, where issues of core-periphery structure are reshaped by the digital nature of the technology and enable, for example, the continuous entry of new organizations (as opposed to the traditional “dominant design and shake-out” kind of dynamics). In other words, the strength of the evolutionary approach is demonstrated by its own evolutionary dynamics, which have generated and enabled new streams of work—even if the new streams’ intellectual debt to evolutionary foundations (e.g., on ecosystems and digital platforms) is not always as explicit as it should be. Table 3 shows this evolution, illustrating it with the specific components of our AI analysis captured by each account.

7. The Value-Add of an Evolutionary Approach

The analysis of the empirical features of “roots” as well as “branches” of evolutionary analysis is more resource- and effort- intensive than what is usually provided in the discourse about AI in economics or policy circles. It requires us to consider the idiosyncrasies and strategies of the key players before we even contemplate national and geopolitical concerns. The question is, what do we gain from this additional level of complexity?

The answer is threefold. The first dimension is phenomenological. As a recent post on towardsdatascience.com lamented, “The AI Ecosystem is a MESS. Why is it impossible to understand what AI companies really do?” (Dhinakaran, 2020). Once we understand the

organization of a new field, drawing on diverse sources as well as primary research, we can draw a map that can guide policy and strategy alike. When, for instance, policymakers say they want to “support AI,” what exactly do they (or should they) mean? Is it to support Big Tech? Their ecosystem? Other specialists? Firms using off-the-shelf AI solutions? Alternative providers of libraries, so as to reduce dependence on Big Tech? Enhance AI use? If so, by what types of firm? Given the asymmetric use and impact of AI, what form does “supporting AI” take, and who exactly stands to benefit from it? On the basis of such an analysis, we can see the implications of well-intentioned but loose policy prescriptions, and tailor our approach accordingly.

The second benefit of the evolutionary approach is epistemological. For better or worse, the careful empirical work done by economists abstracts away from the very phenomena that an evolutionary approach must consider. Nelson & Winter (1982) strongly argued that treating “technology as a residual,” bundling together heterogeneous firms competing as a “production function,” can be misleading, as it neglects the premises that underpin corporate development. For our context, to understand whether or not AI will advance (and whether such an advance will be good or bad), we need to understand the dynamics of who produces AI, how they monetize their advantage, and how this interacts with the attributes and capabilities of those who use AI—as we have aimed to do in this paper.

The third benefit is pragmatic (albeit with theoretical implications). Our evolutionary approach provides a fresh set of responses to existing questions. For instance, it helps us rethink the role of AI as a GPT, and the sense in subsidizing AI; and it helps us revisit whether AI, as speculated by Aghion, Bergeaud, et al (2019), allows for firms with AI investments to expand into different “verticals,” thus transforming the economy. The next two sections explain why the insights based on our analysis (and the resulting prescriptions) *differ* from the established wisdom. We then explain why our analysis is valuable as we seek to understand the interplay of agency and structure and the levers for change—key questions for policy analysis—and conclude with two areas where an evolutionary analysis can help address some important open questions in strategy.

7.1. AI as a GPT, and an Evolutionary Rethink on What This Implies

Significant research has gone into exploring whether AI is, indeed, a GPT (see Cockburn et al, 2019; Goldfarb et al, 2020; Brynjolfsson et al, 2019). The interest in GPTs emanates from their promise to *ultimately* increase productivity and growth—and, more pragmatically, from the fact that GPTs are deemed to merit public funding support in light of their externalities. In contrast, our analysis, rather than focusing on whether AI is “General Purpose” or not, looks

at how AI is produced and used—and, on this basis, also yields different prescriptions on whether public funds should be used to support it.

First, we find that AI was produced incidentally for the need of certain tech companies, and was partly opened up not only to drive AI tech development, but also because those firms might stand to benefit. Big Tech are both the heaviest users of AI and its key beneficiaries; among them, hyperscalers have the additional motivation to provide AI, as this drives the demand for their computing services. As such, the under-provision issue is solved for the large firms—yet only because it makes sense for Big Tech to provide such libraries and promote AI research, since they hold the processing-power bottleneck. This suggests a very different case for potential state involvement—not to encourage any AI production, but rather to stop too much power shifting from states to firms.

Using an evolutionary lens, our attention also shifts from a focus on the aggregate impact of a technology to how particular sectoral and national patterns of innovation drive both its adoption and its implementation. The extent to which AI becomes a priority and its ability to shape productivity is not a function of the underlying technology and complementary investments alone; rather, it is crucially dependent on how particular actors engage and interact. This contrasts with the concern in much of the GPT literature that the lack of coordination between different parties affected by a systemic innovation such as AI will lead to under-provision, and may dampen its impact. We find that country-level and sectoral dynamics are critical in this regard—partly explaining differential AI uptake. Empirically, we find a remarkably close connection between leading academic institutions and AI-producing firms. We also find that the creation of ecosystems that connect AI developers and AI firms have helped mitigate coordination issues, thanks to modularity, and with tools such as AI libraries and developer communities oiling the wheels of innovation.

Our analysis also points to the issues that arise from the heterogeneity of capabilities and incentives to engage with AI. Evolutionary approaches with a deeper appreciation of the nature and sources of heterogeneity can move well beyond size heterogeneity, which seems to be dominant in some economic analyses of AI's heterogeneous impact (see Mihet & Phillipon, 2019; Aghion, Bergeaud et al, 2019). The main issues that we see, drawing on our evolutionary map, is that only a few firms are incented to use and produce AI, and digital sophistication is a precondition for AI use and productivity gains. As such, rather than subsidizing AI across the board, policy may need to address the question of *who engages* in AI. We find that in some countries, big AI players are pushing for AI adoption (e.g., Alibaba in China). This outlines an important challenge for the U.S. and EU as they try to identify AI-friendly policies.

7.2. Does AI Cause Firms to Move to Other Verticals?

Another hypothesis that has been linked to the view of AI as a GPT is that firms that develop competencies in AI may expand into other areas (verticals). Industry surveys also identify this as a crucial topic. According to a recent BCG survey,³³ 78% of companies think that their organization will be prepared to pivot into new businesses because of AI, and 79% are interested in AI because they think new organizations using AI will enter their market. This is a theme that also pervades the recent book by Iansiti & Lakhani (2020), which draws on a number of different settings—in particular, Big Tech players like Microsoft and Alibaba, who have extended into a number of different segments.³⁴ On the other hand, it is not exactly clear whether it is *AI* that (mainly) causes these benefits, or an agile, digital operating model—or, more importantly, whether it is the fact that these firms have access to customer data and relationships. What is clear, however, is that when all these factors coincide, the impact is significant.

This thesis appears to be part of the mainstream vernacular economics, based, among others, on a theoretical argument by Aghion, Bergeaud et al. (2019) that new technologies lower the overhead costs of spanning multiple markets and allow the most productive firms to expand, which would then, in the presence of heterogeneous firms, lead to broader scope and greater profitability disparity. Oddly, the only detailed empirical investigation of AI (Babina et al, 2020: Table 12) does *not* find direct evidence of this thesis, as there is no statistical relation between investment in AI and shifts into other NAICS forms. So how accurate is this belief? Our approach would suggest caution. As we note in Sections 3 and 4, several components of the AI ecosystem (from physical technology infrastructure to AI platforms) are open source or pay-by-use. We have not found evidence of any firm active in AI that can clearly benefit from engaging in all downstream/application activities. As such, there is a risk that the expansion of firms into more verticals has less to do with AI, and more to do with owning information or the customer relationship, or the creation of multi-product ecosystems that can lock customers in (Jacobides & Lianos, 2021). At the technical level, while there may be a few AI models that can be applied to a broad set of phenomena, it is increasingly clear that the understanding of the context cannot be readily separated from the modeling side. This means that for AI

³³ Source: Ransbotham et al, 2020. The sample size was over 3000 firms globally. Respondents to the questionnaire were executives of companies across industrial sectors including aerospace, agriculture, automobile, chemicals, construction and real estate, consumer goods, electronics, entertainment, financial services, health care services, logistics, manufacturing, oil and gas, pharmaceuticals, retail, telecom, transportation and travel, and utilities.

³⁴ Iansiti and Lakhani go into the structural detail, stressing the importance not only of AI investments, but also of a radical rethink of how firms can be structured to take advantage of the opportunities that AI offers—providing some micro-evidence to support many of the observations of Brynjolfsson et al (2019).

production, the existence of domain expertise is an important complement to core AI capabilities. This can be confirmed by the success of firms that specialize in particular domains or industry verticals (e.g., Workday the HR platform; Ping An’s OneConnect financial fraud detection service).³⁵ This may also explain why expansion into new verticals by Big Tech is not met with unequivocal success—exemplified by the Facebook Dating stumble and Uber’s choice to sell off its autonomous car assets.³⁶

7.3. Showcasing the Endogeneity of the Institutional Environment

Our analysis has emphasized the interplay between the technical system (in Section 2) and the institutional structure of production (in Section 3), which helps us understand not only the nature and incentives of participants (and the consequent policy challenges) but also the fact that, as Nelson (1994) argues, there is a co-evolution of technology, sectoral structures, and related institutions. This is made evident by our analysis, in Section 5, of the different trajectories of AI development in the three key contexts. These illustrate the endogeneity of institutions to their environment and serve as a reminder that little is determined by technology alone. Our research shows that the ecosystem-level outcomes we observe are the consequence of strategic choices (which, until now, did not have to contend with much regulatory or even public scrutiny). Technology itself, and the concomitant architecture of rules and roles (Jacobides et al, 2006) has emerged in ways congenial to today’s technology giants on both sides of the Pacific.

We have speculated about the unavoidability of Google-like giants above, with reference to their choice to adopt a marketing business model, in contrast to the founders’ own early ideas. We have argued that their shift was related to the high appropriability of data (and of rents from data), enabled by loose regulations. Yet, technology per se could have supported either an open or a closed solution. This interplay of technology and agency has been at the core of the evolutionary discussion since the very beginning, and as we consider what is distinct about digital strategy (see Adner et al, 2019), it is worth bearing this in mind. AI is not merely “a technology,” just as search is not “a technology.” It is the result of a complex web of choices, mediated by regulatory (in)action—which, along with geopolitical constraints, is about to become crucial (Jacobides, Bruncko, & Langen, 2020).

Our analysis of the way AI is used (in Section 4) also points to end-to-end interdependencies that cut across the different layers of AI enablement, production, and consumption. Users,

³⁵ <https://www.chinadaily.com.cn/a/202008/24/WS5f431339a310834817262255.html>

³⁶ See <https://www.bostonglobe.com/lifestyle/2019/11/08/facebook-new-dating-service-flopping-tried-for-week-find-out-why/LT6j2wsvMepzTA9Y19CqqJ/story.html> and <https://www.economist.com/business/2020/12/10/why-is-uber-selling-its-autonomous-vehicle-division>

relying on AI-enabled applications, contribute data to the ecosystem that allow for the continuous improvement and refinement of the algorithms on which those applications build. And, in so doing, they reinforce the dominant position of a few technology giants. These feedback loops are continuous and ubiquitous, generating substantial concerns in terms of who benefit really from them—e.g., Zuboff (2015). While the role of users in improving technologies at the point of application is not new (see e.g., Nuvolari, 2005 for a broad historical excursus), the seamless, digital connectedness enabled by AI-as-technology is, in our view, a unique feature that we have yet to fully grasp in terms of its economic, managerial, and even psychological implications. This sets in motion powerful economic forces with which society, polity, and regulators will have to contend, obliging them to update their playbook accordingly (Jacobides & Lianos, 2021).

7.4. Revisiting Strategic Choices Ahead

Our analysis can also help us reconsider the key strategy and policy dilemmas that we face. First, as Edge computing becomes more prevalent (i.e., carried out on smaller, local devices such as cameras and phones), chip manufacturers like Nvidia and device manufacturers like Huawei herald their AI-compliant devices and their new chip architectures, which make Edge functions more effective. The question becomes how current leading Cloud providers of AI react and adapt in the face of the rising importance of Edge computing. This typifies the contemporary strategic challenges of industry convergence and competition that comes from firms rooted in different environments (Jacobides, 2010; Kim et al, 2015).³⁷ The implication is that firms such as telcos, supported by their suppliers, are trying to compete with hyperscalers for part of the value-add that AI can provide. Regulation may play an important part in these struggles. Initiatives such as GDPR/ePrivacy, PIS20, and the EU's Digital Services / Market Act will determine the attractiveness of each business model (see Jacobides, Bruncko & Langen, 2020). The extent to which Edge-enabled applications (that collect and process data in a decentralized way) are a substitute or a complement to top-down Cloud-based computing depends on the strategic design of interoperability standards.³⁸ At this

³⁷ For Cloud providers, the key future focus areas appear to be both training intelligence on the Cloud and continuing to control data end-to-end. Edge firms (connectivity service providers, hardware OEMs, established industrial goods players, and CDNs) better enable the Cloud by sending data for storage or inference. An example of these “architectural” strategic battles can be seen through the development of Multi-Access Edge Computing (i.e., servers connected near 5G towers), which would allow connectivity service providers like telecommunications firms to enable AI applications.

³⁸ One can think of technical solutions that would be compatible with both Google's Cloud technologies and Huawei's Edge solutions. Currently, however, none are, due to former President Trump's decision to wage a geopolitical war against Huawei.

stage, Edge and Cloud solutions rely on the same few large players end-to-end. Yet regulators might decide to push for greater standard homogenization and interoperability.

Beyond such rarefied architectural battles, regular firms *using* AI technologies may benefit from considering what evolutionary approaches have taught us. AI, like other new technologies, requires “absorptive capacity” (Cohen & Levinthal, 1990), so that success requires a base understanding of AI in order to benefit from it. By analogy to Brusoni et al (2001), firms need to “know more than they do” to be able to effectively respond to AI. Interestingly, though, AI is enabling a few giant firms to know more and more, thanks to what downstream firms and users do. The role of knowledge integration (Jacobides et al. 2009) or dynamic architectural capabilities (Baldwin, 2018) is likely to become more and more salient to explain competitiveness and growth in the AI ecosystem.³⁹ Several organizational skills need to complement AI (Brynjolfsson et al, 2019; Ransbotham et al, 2020; Sudarshan et al, 2018), and these are likely to differ between AI infrastructure / enablers and application providers, between AI producers and those who merely consume it. As such, we hope that the map that we provide will help better explain and prescribe.

Our evolutionary approach, with its emphasis on inter-firm heterogeneity and its evolution, can also fruitfully combine with the recent empirical interest in the emergence of “superstar firms” that expand broadly, and grow in scope (Lashkari et al., 2018; Autor et al., 2020). This can also inform our understanding of competition law, which has begun to grapple with the thorny area of antitrust and may need to be broadened still further to consider ecosystem dynamics. This will become ever more relevant as the geopolitical confrontations between the U.S., the EU, and China may further reshape the landscape, posing challenges for firms, policymakers, and societies. Such problems are both urgent and complex, and we hope that an evolutionary approach, like the one we propose in this article, will prove useful in overcoming them.

³⁹ In an ecosystem where open and owned data, communities of freelancers and employees, Big Tech and start-ups all coexist, the relationship between “doing” and “knowing” is mediated by a complex web of heterogeneous institutions and norms to which we need to give more attention.

References

- Adner R (2017) Ecosystem as structure: An actionable construct for strategy. *J. Manag. Stud.* 43(1): 39–58.
- Adner, R., & Kapoor, R. (2010). Value creation in innovation ecosystems: How the structure of technological interdependence affects firm performance in new technology generations. *Strateg. Manag. J.* 31(3): 306-333.
- Adner R, Puranam P, Zhu F (2019) What is different about digital strategy? From quantitative to qualitative change. *Strategy Sci.* 4(4): 253-261.
- Aghion P, Bergeaud A, Boppart T, Klenow PJ, Li H (2019) A theory of falling growth and rising rents. *NBER Working Paper w26448*.
- Aghion P, Jones B, Jones C (2019) Artificial intelligence and economic growth. Agrawal A, Gans J, Goldfarb A, eds. *The Economics of Artificial Intelligence: An Agenda*. (University of Chicago Press, Chicago, Illinois), 237-290.
- Agrawal A, Gans J, Goldfarb A (2019) *The Economics of Artificial Intelligence an Agenda* (The University of Chicago Press, Chicago).
- Alekseeva L, Azar J, Gine M, Samila S, Taska B (2019) The demand for AI skills in the labor market. *SSRN Electronic Journal*, <https://doi.org/10.2139/ssrn.3470610>
- Allen B, Gish R, Dreyer K (2019) The role of an artificial intelligence ecosystem in radiology. Ranschaert ER, Morozov S, Algra PR, eds. *Artificial Intelligence in Medical Imaging* (Springer, New York), 291-327.
- Arenal, A., Armuña, C., Feijoo, C., Ramos, S., Xu, Z., & Moreno, A. (2020). Innovation ecosystems theory revisited: The case of artificial intelligence in China. *Telecomm. Policy*, 44(6): 101960.
- Arora A, Belenzon S, Pataconi A (2018) The decline of science in corporate R&D. *Strateg. Manag. J.* 39(1): 3-32.
- Autor D, Dorn D, Katz LF, Patterson C, Van Reenen J (2020) The fall of the labor share and the rise of superstar firms. *Q. J. Econ.* 135(2): 645-709.
- Babina T, Fedyk A, He A, Hodson J (2020) Artificial intelligence, firm growth, and industry concentration. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3651052>
- Baldwin C (2018) Bottlenecks, modules, and dynamic architectural capabilities. Teece DJ, Heaton S, eds. *The Oxford Handbook of Dynamic Capabilities* (Oxford University Press, Oxford, U.K.).
- Baldwin C, Clark K (2000) *Design Rules Volume 1: The Power of Modularity* (M.I.T. Press, Cambridge, MA)
- Baldwin C, Woodard C (2009) The Architecture of Platforms: A Unified View. Gawer A, eds. *Platforms, Markets and Innovation* (Edward Elgar Publishing, Cheltenham, U. K.), 19-44.
- Barney J (1991) Firm resources and sustained competitive advantage. *J. Manag. Stud.* 17(1): 99-120.
- Bloom N, Van Reenen J (2010). Why do management practices differ across firms and countries? *J Econ Perspect.* 24(1): 203-24.
- Boudreau K, 2010, Open platform strategies and innovation: Granting access vs. devolving control, *Management science* 56 (10): 1849-1872
- Breschi S, Orsenigo L, Malerba F (2000) Technological Regimes and Schumpeterian Patterns of Innovation. *Econ. J.* 110(463): 388-410.
- Bresnahan T, Trajtenberg M (1995) General-purpose technologies ‘Engines of growth’? *J. Econom.* 65(1): 83-108.
- Brock J, Wangenheim F (2019) Demystifying AI: What digital transformation leaders can teach you about realistic artificial intelligence. *Calif. Manag. Rev.* 61(4): 110-134.

- Brin S, Page L (1998) The anatomy of a large-scale hypertextual web search engine. *Comput. Netw. ISDN Sys.* 30(1-7): 107-117.
- Brusoni S, Prencipe A, Pavitt K (2001) Knowledge specialization, Organizational coupling, and the boundaries of the firm: Why do firms know more than they make? *Adm. Sci. Q.* 46(4): 597-621.
- Brynjolfsson E, Rock D, Syverson C (2019) Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. Agrawal A, Gans J, Goldfarb A, eds. *The Economics of Artificial Intelligence: An Agenda*. (University of Chicago Press, Chicago, Illinois), 23-60.
- Brynjolfsson E, Hitt L (1998) Beyond the IT productivity paradox. *Commun. ACM.* 41(8): 49-55.
- Bughin J, Seong J, Manyika J, Chui M, Joshi R (2019) *Notes from the AI frontier: Modeling the impact of AI on the world economy*. Retrieved [from here](#) February 11, 2021
- Burström T, Parida V, Lahti T, Wincent J (2021) AI-enabled business-model innovation and transformation in industrial ecosystems: A framework, model and outline for further research. *J. Bus. Res.* 127: 85-95.
- Candelon F, Reichert T, Duranton S, Charme di Carlo R, Stokol G (2020) *Deploying AI to Maximize Revenue*. Boston Consulting Group. Retrieved [from here](#) February 11, 2021
- Cefis E, Ciccarelli M (2005) Profit differentials and innovation. *Econ. Innov. New Technol.* 14(1-2): 43-61.
- Chan-Olmsted SM (2019) A review of artificial intelligence adoptions in the media industry. *Int. J. Media Manag.* 21(3-4): 193-215.
- Clough DR, Wu A (2021) Artificial intelligence, Data-driven learning, and the decentralized structure of platform ecosystems. *Acad Manage Rev.* in press.
- Coase RH (1937) Some notes on monopoly price. *Rev. Econ. Stud.* 5(1): 17-31.
- Cockburn I, Henderson R, Stern S (2019) The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis. Agrawal A, Gans J, Goldfarb A, eds. *The Economics of Artificial Intelligence: An Agenda* (University of Chicago Press, Chicago, Illinois), 115–148.
- Cohen W, Levinthal D (1990) Absorptive Capacity: A New Perspective on Learning and Innovation. *Adm. Sci. Q.* 35(1): 128-152.
- David PA (1992) Heroes, herds and hysteresis in technological history: Thomas Edison and ‘The Battle of the Systems’ reconsidered. *Ind. Corp. Chang.* 1(1): 129-180.
- Dhar P (2020) The carbon impact of artificial intelligence. *Nat. Mach. Intell.* 2: 423–425.
- Dhinakaran A (2020) *The AI ecosystem is a mess*. Retrieved February 11, 2021, from <https://towardsdatascience.com/the-ai-ecosystem-is-a-mess-c46bdfbf43e4>
- Economist (2017) *Data is giving rise to a new economy*. Retrieved February 12, 2021, from <https://www.economist.com/briefing/2017/05/06/data-is-giving-rise-to-a-new-economy>
- Etchemendy J, Li F (2020) *National research Cloud: Ensuring the continuation of American innovation*. Retrieved February 11, 2021, from <https://hai.stanford.edu/blog/national-research-cloud-ensuring-continuation-american-innovation>
- Ethiraj S (2007) Allocation of inventive effort in complex product systems. *Strateg. Manag. J.* 28(6): 563-584.
- Etzkowitz H, Leydesdorff L (1995) The Triple Helix - University - Industry - Government Relations: A Laboratory for Knowledge Based Economic Development. *EASST Review*, 14(1): 14-19.
- European Commission (2020) *White paper on Artificial intelligence: A European approach to excellence and trust*.
- Executive Office of the President (2019) *Maintaining American leadership in artificial intelligence*. Executive Order 13859.

- Freeman C, Louçã F (2001) *As Time Goes By: From the Industrial Revolutions to the Information Revolution* (Oxford University Press, Oxford, U. K.).
- Furman, J. L., & Teodoridis, F. (2020). Automation, research technology, and researchers' trajectories: Evidence from computer science and electrical engineering. *Organ. Sci.* 31(2): 330-354.
- Furman J, Seamans R (2019) AI and the economy. *Innov. Policy Econ.* 19(1): 161-191.
- Garbuio M, Lin N (2019) Artificial intelligence as a growth engine for health care startups: Emerging business models. *Calif. Manag. Rev.* 61(2): 59-83.
- Gawer, A., & Cusumano, MA (2002). *Platform leadership: How Intel, Microsoft, and Cisco drive industry innovation* (Boston, MA: Harvard Business School Press).
- Gerbert P, Hecker M, Steinhäuser S, Ruwolt P (2017) *Putting Artificial Intelligence To Work*. Retrieved February 12, 2021, from <https://www.bcg.com/fr-fr/publications/2017/technology-digital-strategy-putting-artificial-intelligence-work>
- Goldfarb A, Taska B, Teodoridis F (2020). Artificial Intelligence in Health Care? Evidence from Online Job Postings. *AEA Papers and Proceedings.* 110: 400-404.
- Gawer A, Cusumano MA (2014) Industry platforms and ecosystem innovation. *J. Prod. Innov. Manag.* 31(3): 417-433.
- He X, Zhao K, Chu X (2020) AutoML: A Survey of the State-of-the-Art. *Knowl Based Syst.* 212: 106622.
- Herweijer C, Combes B, Ramchandani P, Sidhu J (2018) *PwC, World Economic Forum, Stanford Woods Institute for the Environment. Harnessing Artificial Intelligence for the Earth.* (World Economic Forum, Geneva)
- Hobday, M. (1998). Product complexity, innovation and industrial organisation. *Research Pol.* 26(6): 689-710.
- Hughes TP (1993) *Networks of power: electrification in Western society, 1880-1930* (Johns Hopkins University Press, Baltimore, Maryland)
- Iansiti M, Lakhani KR (2020) *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world* (Harvard Business School Press, Boston).
- Jacobides MG (2010) Strategy tools for a shifting landscape. *Harv. Bus. Rev.* 88(1): 76–85.
- Jacobides MG, Brusoni S, Prencipe A (2009) Strategic dynamics in industry architectures and the challenges of knowledge integration. *Eur. Manag. Rev.* 6(4): 209-216.
- Jacobides MG, Bruncko M, Langen R (2020) Regulating Big Tech in Europe: Why, so what, and how understanding their business models and ecosystems can make a difference. *Evolution Ltd White Paper*, available at <https://www.evolutionltd.net/post/regulating-big-tech-in-europe>
- Jacobides MG, Cennamo C, Gawer A (2018) Toward a theory of business ecosystems. *Strateg. Manag. J.* 39(8): 2255-2276.
- Jacobides MG, Knudsen T, Augier M (2006) Benefiting from innovation: Value creation, value appropriation and the role of industry architectures. *Res. Policy.* 35(8): 1200-1221.
- Jacobides MG, Lianos I (2021) *Ecosystems and competition law in theory and practice.* Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3772366
- Jacobides MG, MacDuffie JP, Tae CW (2016) Agency, structure, and OEM dominance: Change and stability in the automotive sector. *Strateg. Manag. J.* 37(9): 1942-1967.
- Jacobides MG, Tae CJ (2015) Kingpins, Bottlenecks, and Value Dynamics Along a Sector. *Organ. Sci.* 26(3): 889-907.
- Jacobides MG, Winter SG (2005) The co-evolution of capabilities and transaction costs: Explaining the institutional structure of production. *Strateg. Manag. J.* 26(5): 395-413.

- Jacobides MG, Winter SG (2012) Capabilities: Structure, Agency, and Evolution. *Organ. Sci.* 23(5): 1365-1381.
- Kaplan J, McCandlish S, Henighan T, Brown TB, Chess B, Child R, Gray S, Radford A, Wu J, Amodei D (2020) *Scaling Laws for Neural Language Models*. arXiv preprint arXiv:2001.08361.
- Kim N, Lee H, Kim W, Lee H, Suh JH (2015) Dynamic patterns of industry convergence: Evidence from a large amount of unstructured data. *Res. Policy.* 44(9):1734-1748.
- Kumar V, Rajan B, Venkatesan R, Lecinski J (2019) Understanding the role of artificial intelligence in personalized engagement marketing. *Calif. Manag. Rev.* 61(4): 135-155.
- Lashkari, D., Bauer, A., & Boussard, J. (2018). Information technology and returns to scale. Available at SSRN 3458604. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3458604
- Lee KF (2018). *AI Superpowers: China, Silicon Valley, and The New World Order* (Houghton Mifflin Harcourt, Boston, MA).
- Lundvall BÅ (1992) *National Systems of Innovation: Towards a Theory of Innovation and Interactive Learning* (Anthem Press, New York).
- Lundvall, BÅ (2007) National innovation systems—Analytical concept and development tool. *Ind Innov.* 14(1): 95-119.
- Malerba F (2004) Sectoral systems of innovation: basic concepts. Malerba F, eds. *Sectoral Systems of Innovation Concepts, Issues and Analyses of Six Major Sectors in Europe* (Cambridge University Press, Cambridge, U.K.), 9-41.
- Malerba F, Nelson R, Orsenigo L, Winter S (2016) *Innovation and the Evolution of Industries: History-Friendly Models* (Cambridge University Press, Cambridge, U. K.)
- Malerba F, Orsenigo L (1999) Technological entry, exit and survival: an empirical analysis of patent data, *Res. Policy.* 28(6): 643-660.
- Marcus G (2018) *Deep Learning: A Critical Appraisal*. ArXiv: abs/1801.00631.
- Metcalf J (1994) Evolutionary Economics and Technology Policy. *Econ. J.* 104(425): 931-944.
- Mihet R, Philippon T (2019) The Economics of big data and artificial intelligence. Choi JJ, Ozkan B, eds. *Disruptive Innovation in Business and Finance in the Digital World* (Emerald, Bingley, U. K.) 29-43.
- Miller R, Laird R (2019) *Reshaping the global it landscape: The impact of hyperscale data centers*. Retrieved February 11, 2021, from <https://datacenterfrontier.com/reshaping-the-global-it-landscape-the-impact-of-hyperscale-data-centers/>
- MIT (2017) *Three-Year Action Plan for Promoting Development of a New Generation Artificial Intelligence Industry (2018–2020)*. Ministry of Industry and Information Technology of China.
- Mou XM (2019) *Artificial intelligence: Investment trends and selected industry uses*. International Finance Corporation.
- Nelson, R R (1994). The co-evolution of technology, industrial structure, and supporting institutions. *Ind. Corp. Chang.* 3(1): 47-63.
- Nelson RR (2020) A perspective on the evolution of evolutionary economics. *Ind. Corp. Chang.* 29(5): 101–1118.
- Nelson RR, Winter SG (1982) *An Evolutionary Theory of Economic Change*. (Belknap Press, Massachusetts)
- Nuvolari A (2005) Open source software development: Some historical perspectives. *First Monday*, 10(10): 3.
- Parker, G. G., Van Alstyne, M. W., & Choudary, S. P. (2016). *Platform revolution: How networked markets are transforming the economy and how to make them work for you* (WW Norton & Company).

- Pisano GP, Teece DJ (2007) How to Capture Value from Innovation: Shaping Intellectual Property and Industry Architecture. *Calif. Manag. Rev.* 50(1): 278–296.
- Pratt GA (2015) Is a Cambrian Explosion Coming for Robotics? *J Econ Perspect.* 29(3): 51-60.
- Ransbotham S, Khodabandeh S, Kiron D, Candelon F, Chu M, LaFountain B (2020) *Expanding AI's Impact with Organizational Learning*. MIT Sloan Management Review and Boston Consulting Group.
- Rock D (2019) Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence. *SSRN Electronic Journal*.
- Rosenberg N (1982) *Inside the Black Box: Technology and Economics*. (Cambridge University Press, Cambridge, U.K.)
- Saviotti PP (2005) Innovation systems and evolutionary theories. Edquist C, eds. *Systems of Innovation: Technologies, Institutions, and Organizations* (Pinter Publishers/Cassell Academic, London, U. K.).
- Schumpeter JA (1942) *Capitalism, Socialism and Democracy*.
- Schumpeter JA (1912) *Theorie der wirtschaftlichen Entwicklung* (Duncker & Humblot, Leipzig)
- Sharir O, Peleg B, Shoham Y (2020) *The cost of training NLP models: A concise overview*. arXiv preprint arXiv:2004.08900v1.
- Shrestha YR, Ben-Menahem SM, Von Krogh G (2019). Organizational decision-making structures in the age of artificial intelligence. *Calif. Manag. Rev.* 61(4): 66-83.
- Simon JP (2019) Artificial intelligence: scope, players, markets and geography. *Digit. Policy, Regul. Gov.* 21(3): 208-237.
- Sosa LM (2013). Decoupling market incumbency from organizational prehistory: Locating the real sources of competitive advantage in R&D for radical innovation. *Strateg. Manag. J.* 34(2): 245-255.
- Simon HA (1970) *The Sciences of the Artificial* (M.I.T. Press, Cambridge, Massachusetts).
- Sudarshan PN, Abhishek V, Gupta G (2018) *Artificial Intelligence: Why businesses need to pay attention to artificial intelligence?* Retrieved [from here](#) February 12, 2021.
- Tambe, P., Hitt, L. M., Rock, D., & Brynjolfsson, E. (2019). IT, AI and the Growth of Intangible Capital. Available at SSRN 3416289.
- Teece D, Pisano G, Schuen A (1997) Dynamic capabilities and strategic management. *Strateg. Manag. J.* 18(7): 509-533.
- Varian HR (2019) Artificial Intelligence, Economics, and Industrial Organization. Agrawal A, Gans J, Goldfarb A, eds. *The Economics of Artificial Intelligence* (University of Chicago Press, Chicago), 399-422.
- Winter SG (1995) Four Rs of profitability: rents, resources, routines, and replication. Montgomery CA, eds. *Resource-based and Evolutionary Theories of the Firm: Towards a Synthesis* (Springer, Boston, MA), 147-178.
- World Economic Forum (2018) *The Future of Jobs Report 2018*. Geneva, World Economic Forum. Retrieved February 12, 2021 from <https://www.weforum.org/reports/the-future-of-jobs-report-2018>
- Zigbee Alliance (2020) *Amazon, Apple, Google, and the Zigbee alliance and its board Members FORM industry working group to develop a new, open standard for smart home device connectivity*. Retrieved February 11, 2021, from https://zigbeealliance.org/news_and_articles/connectedhomeip/
- Zuboff S (2015) Big other: surveillance capitalism and the prospects of an information civilization. *J. Inf. Technol.* 30(1): 75-89.

Table 1: AI Production and Consumption Landscape

		AI Consumption	
		Sell AI-production services for clients to consume it, but do not use internally	Internal consumption to enhance offer or company’s performance
AI Production	Production mostly in-house	AI Giants: end-to-end integration across the AI stacks, they have the capability to produce the AI necessary for internal and external use (e.g., Google, Amazon, Microsoft, Alibaba, Tencent, Baidu)	
	Balanced mix of in-house and purchased production	AI Creators: have the capability to produce or customize some of the AI sold to their clients, but rely on AI giants for “AI basics” (e.g., Accern, MonkeyLearn, Leivity, AI consulting services)	AI-powered Operators: leverage AI in their day-to-day operations and their offer, using both AI Giants services and internal capability to produce the AI necessary for critical functions/operations (e.g., Facebook, Uber, Spotify, ByteDance, traditional companies with internal AI Powerhouse - Walmart, Ping An, ...)
	AI production mostly purchased	AI Traders/Integrators: purchase and sell off-the-shelf AI solutions or use-cases, adding commercial and marketing efforts (e.g., bundling, repackaging, branding) or supporting integration with client ecosystem, but without AI improvements (e.g., translation companies using Google Translate services, CRM consultants integrating Salesforce or MS offers, ...)	AI Takers: leverage only or mostly off the shelf solutions to enable critical business functions (e.g., digital natives outsourcing most of their AI production, traditional companies with limited internal AI capabilities)

Source: BCG primary and secondary research; AI implementation projects including multiple interviews

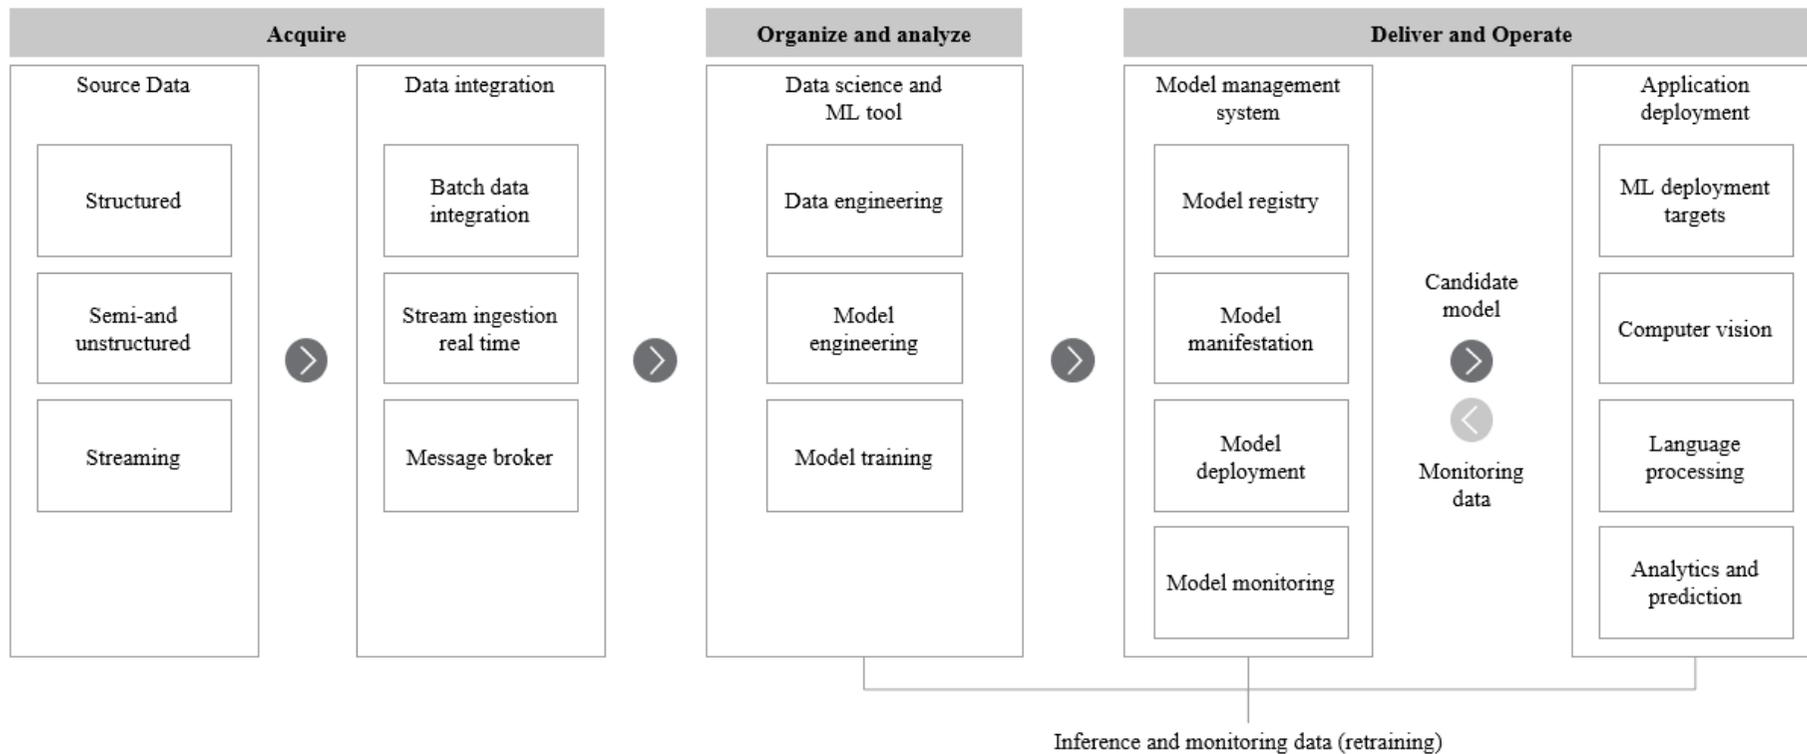
Table 2: Attitudes towards AI in China, the EU, and U.S./Canada

Question	China	U.S. and Canada	EU
End users of AI generally trust the AI solution's decisions	86%	39%	45%
End users of AI believe that AI improve business outcomes	88%	54%	62%
End users of AI understand the reasons behind the specific recommendations	85%	39%	42%
End users of AI understand the limitations of the AI system	80%	35%	28%
We are interested in AI because Customers will ask for AI-driven offerings	87%	61%	63%
We are interested in AI because Suppliers will offer AI-driven products and services	89%	72%	69%
We understand the cost of developing AI	90%	50%	54%
We set up cross-functional teams from the beginning	84%	56%	61%

Source: BCG/MIT Survey, 2020. The sample size was over 297 firms globally (with 114 firms in China, 128 in the U.S./Canada, and 65 in EU. Firms surveyed had revenues over \$500 million in the past fiscal year. Respondents to the questionnaire were executives of companies across industrial sectors including aerospace, agriculture, automobile, chemicals, construction and real estate, consumer goods, electronics, entertainment, financial services, health care services, logistics, manufacturing, oil and gas, pharmaceuticals, retail, telecom, transportation and travel, and utilities.

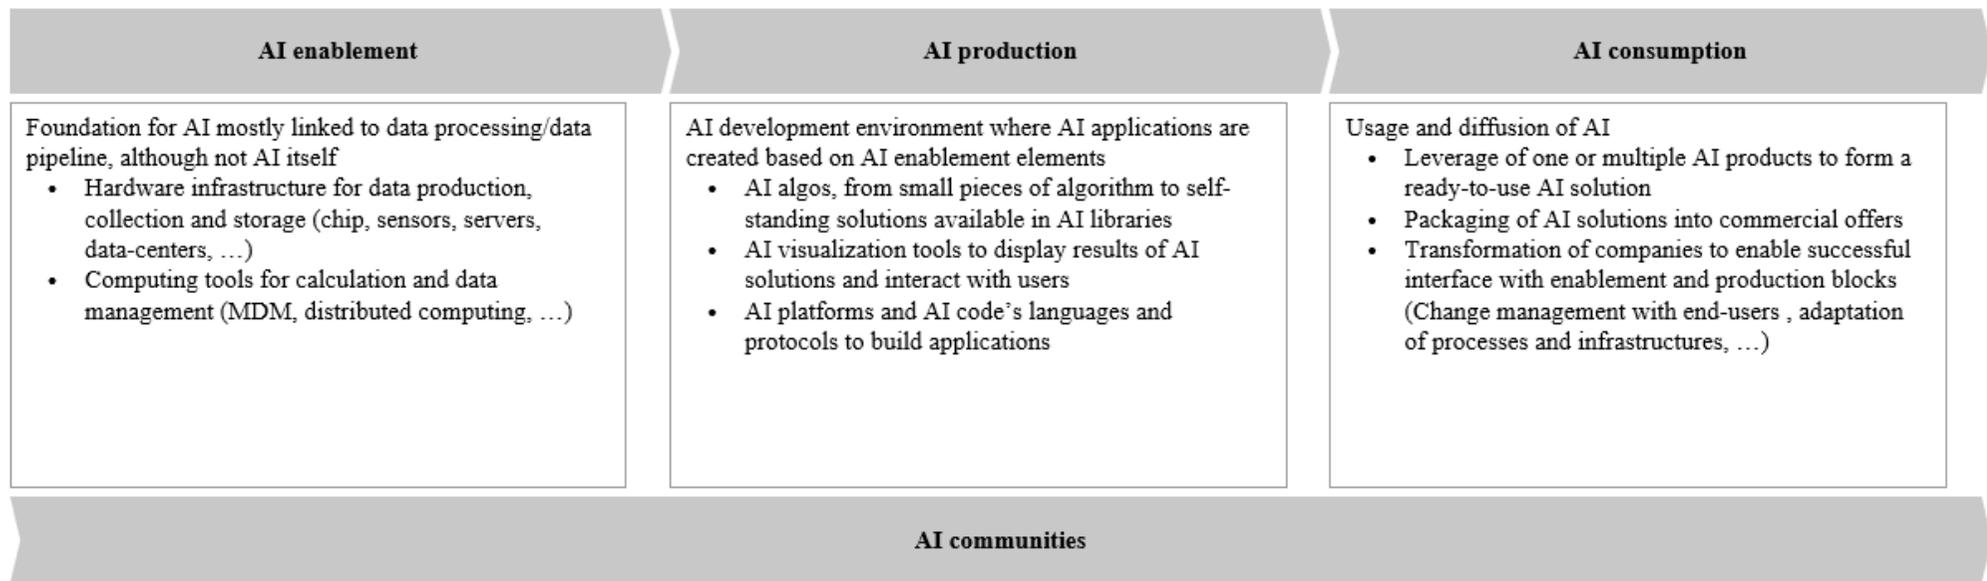
Table 3: The Evolution of Evolutionary Typologies of Systems and the world of AI

	Foundation → Extension		Foundation → Extension		Foundation → Extension	
	National systems of innovation	Triple helix model	Sectoral systems of innovation / Industry Architecture	Business and innovation ecosystems	Complex technical systems	Digital platforms
Focus	Role of national governments in shaping innovation trajectories, through e.g. regulations and government-sponsored research	Dynamic, open-ended interplay between actors in industry, university systems and government	Analysis of the dynamic interactions between technologies, actors, and institutions in shaping sector-specific patterns / Endogenously evolving rules and roles in the division of labor in a sector	Emphasis on complementors and dynamics of competition and cooperation; direct and indirect network effects; winner-take-all dynamics and innovation around key orchestrators	Attention to the internal dynamics of complex systems, their evolving architectures and interdependencies; emergence and stability of core-periphery structures at the industrial level	Digitally enabled evolvability and generativity of technical systems; continuous entry and persistent emergence of new applications; strategic choices of platform owners to attract and exploit complementors
Indicative contributors	Freeman; Nelson; Lundvall	Etzkowitz & Leydesdorff	Malerba & Orsenigo/ Jacobides & Winter	Adner & Kapoor; Jacobides et al	Hughes; Hobday; Rosenberg	Gawer & Cusumano; Parker et al; Boudreau
AI-related core elements	Competition between Chinese and Western innovation systems, e.g. different approaches to organization of libraries	Increasing role of private business in fundamental AI research; data policy shaping what firms can do with AI; AI state support dynamics	Role of tech giants in developing upstream AI capabilities and the overall data infrastructure; role of coder communities	Emerging role of application-specific, downstream complementors and novel data-driven business models; firm-specific ecosystem strategies	Architecture of the AI system (ie, enablement, production, consumption); connections between firm and technological choices	Strategies of key players and their complementors; (varying) role of data and libraries; (lack of) interoperability between Cloud and Edge technologies



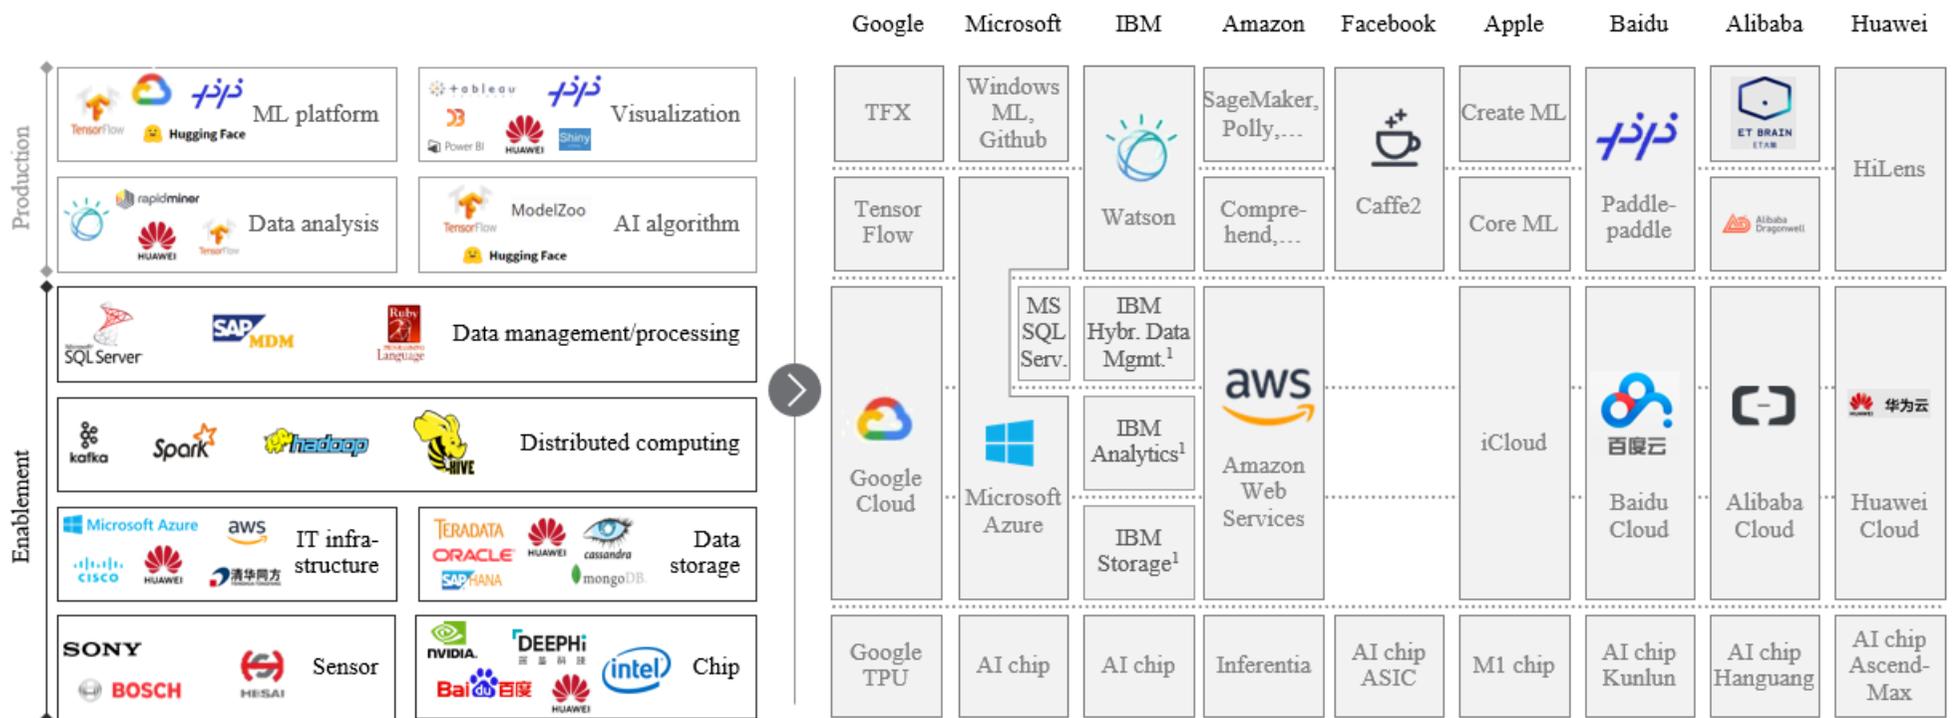
Source: BCG primary and secondary research; AI implementation projects including multiple interviews

Figure 1: The Machine Learning Technical Architecture



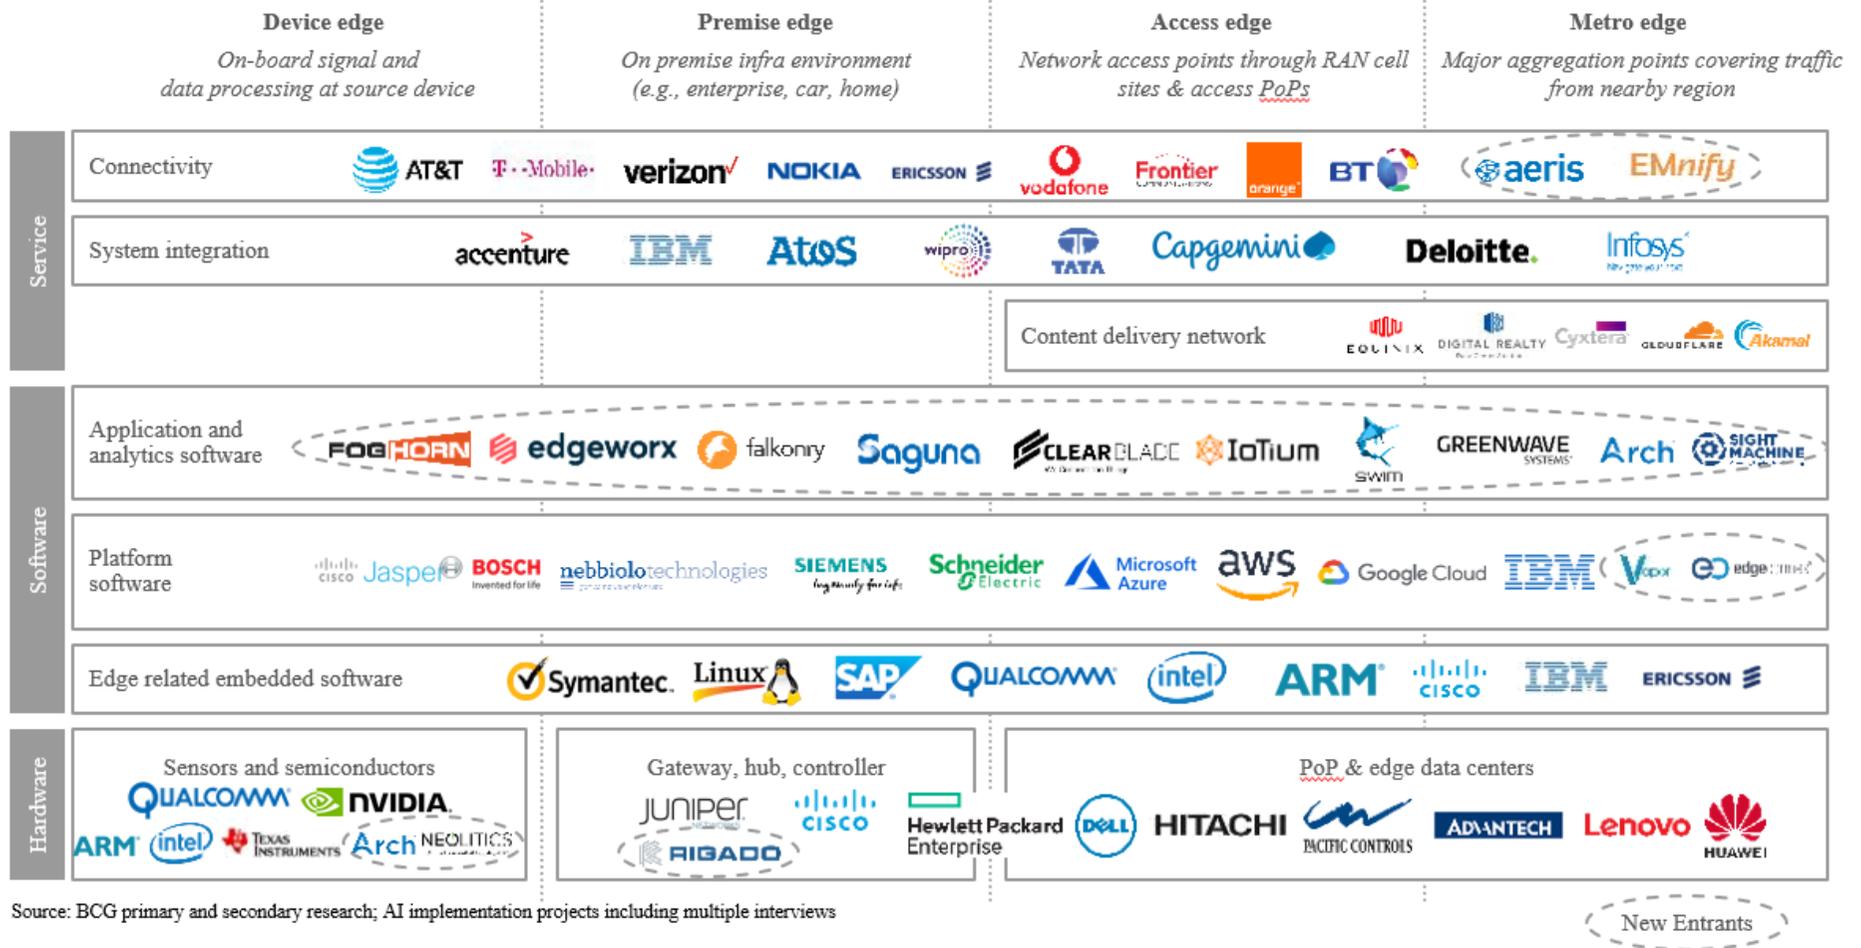
Source: BCG primary and secondary research; AI implementation projects including multiple interviews

Figure 2: AI Development in Three Stages



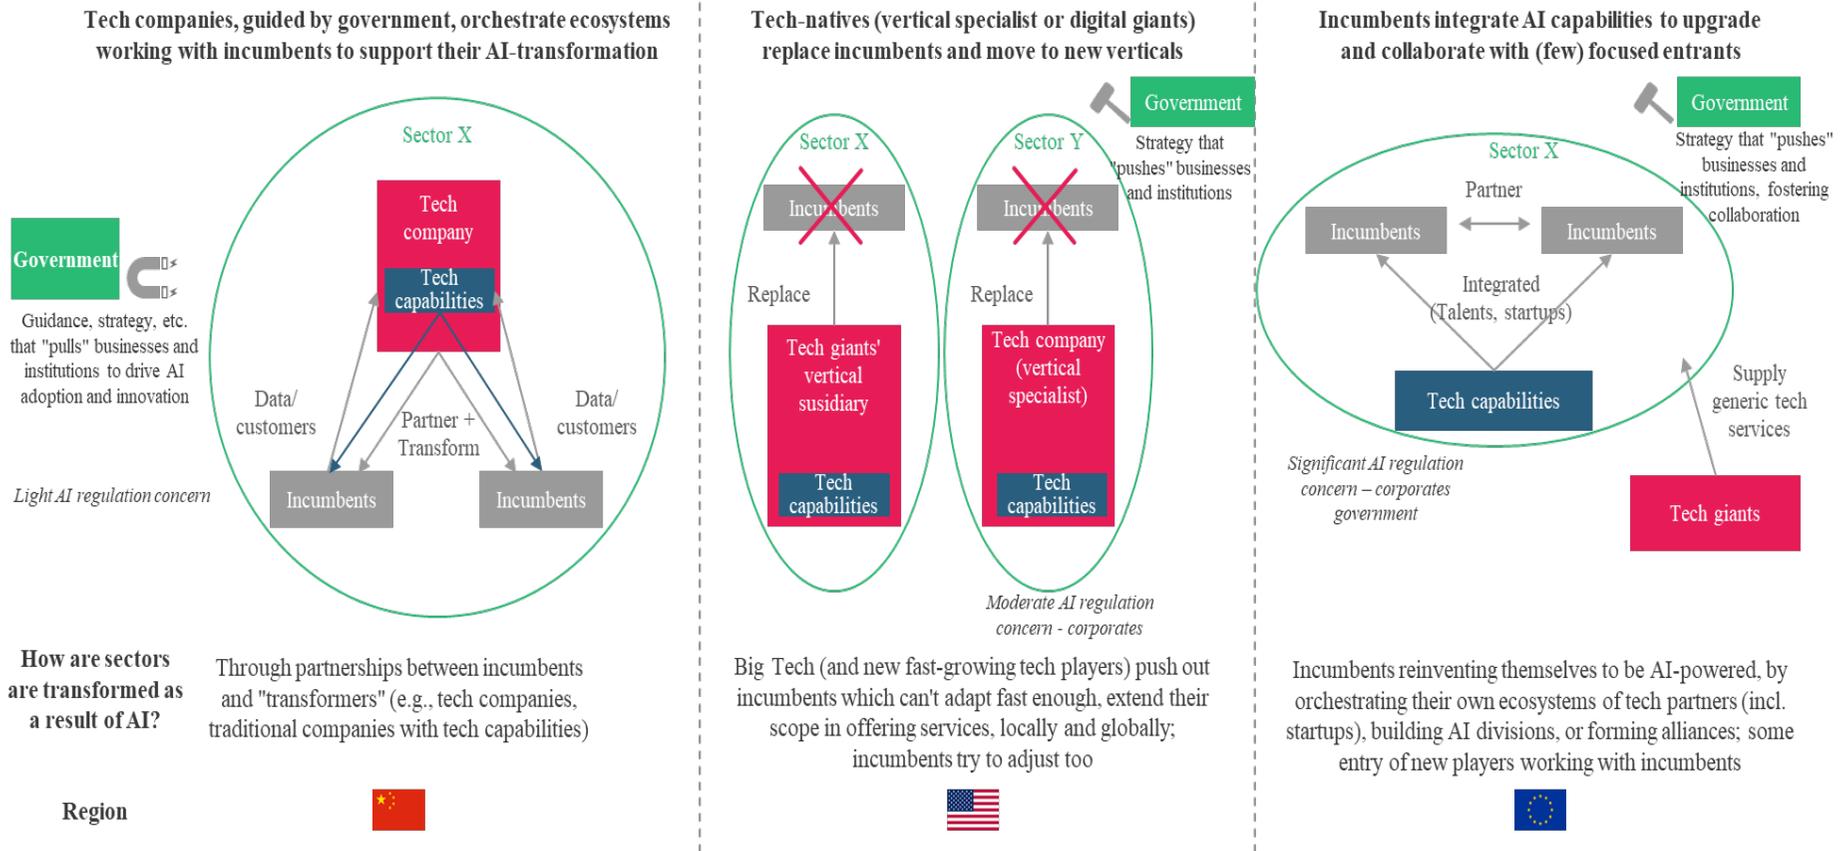
1. Also part of IBM Cloud
 Source: BCG primary and secondary research; AI implementation projects including multiple interviews

Figure 3: AI Technology and IT Stacks: Zoom on AI Enablement and Production Layers and Players



Source: BCG primary and secondary research; AI implementation projects including multiple interviews

Figure 4: Incumbents (*de alio players*) and New Entrants (*de novo players*) in the Growing Edge Market



Note: Details on sources and aspects of how each different country / country block differs are detailed in Online Appendix 3

Figure 5: The differences in the AI sectoral and innovation systems in China, the U.S., and the EU