SCIENCE POLICY RESEARCH UNIT

# SPRU Working Paper Series

SWPS 2021-02 (May)

## Artificial Intelligence's New Clothes? From General Purpose Technology to Large Technical System

Simone Vannuccini and Ekaterina Prytkova

Energy Policy

Sustainable Development

Science and Technology Policy

Innovation and Project Management

Economics of Innovation

UNIVERSITY OF SUSSEX

BUSINESS SCHOOL

# SPRU Working Paper Series (ISSN 2057-6668)

The SPRU Working Paper Series aims to accelerate the public availability of the research undertaken by SPRU-associated people, and other research that is of considerable interest within SPRU, providing access to early copies of SPRU research.



# Guidelines for authors

Papers should be submitted to swps@sussex.ac.uk as a PDF or Word file. The first page should include: title, abstract, keywords, and authors' names and affiliations. The paper will be considered for publication by an Associate Editor, who may ask two referees to provide a light review. We aim to send referee reports within three weeks from submission. Authors may be requested to submit a revised version of the paper with a reply to the referees' comments to swps@sussex.ac.uk. The Editors make the final decision on the inclusion of the paper in the series. When submitting, the authors should indicate if the paper has already undergone peer-review (in other series, journals, or books), in which case the Editors may decide to skip the review process. Once the paper is included in the SWPS, the authors maintain the copyright.

Websites

UoS: **www.sussex.ac.uk/spru/research/swps**
SSRN: **www.ssrn.com/link/SPRU-RES.html**
IDEAS: **ideas.repec.org/s/sru/ssewps.html**

# Artificial Intelligence's New Clothes? From General Purpose Technology to Large Technical System

Simone Vannuccini[1] and Ekaterina Prytkova[2]

[1]Science Policy Research Unit, University of Sussex Business School, University of Sussex
[2]Friedrich Schiller University Jena, Department of Economics and Business Administration
[1]S.Vannuccini@sussex.ac.uk
[2]e.prytkova@uni-jena.de

April 7, 2021

**Abstract**

Artificial Intelligence (AI) has been quickly labelled a General Purpose Technology (GPT) for its many uses and the high expectations built around a technology that can perform tasks associated with natural intelligence. However, for now, the claim "AI equals GPT" is premature, and eventually, taking into account potential future scenarios, it can turn out to be incorrect. In fact, though every GPT is an influential technology, not every influential technology is a GPT. Checking AI against the definitional criteria of GPT, we come to the conclusion that GPT is a misspecified model of AI: what was meant to be a concept for an individual technology in this case is stretched to cover a growing infrastructural, system technology. For example, the pervasiveness featured in the GPT concept seems to be qualitatively different from the largeness that modern AI demonstrates. In this paper, we suggest an alternative framework drawn from the literature on Large Technical Systems (LTS) as more fit to represent the nature of AI. We map the building blocks of LTS on AI and describe its state–of–the–art through this novel viewpoint. This is a timely exercise, as we witness the formation of an AI industry. A correct understanding of its core technology is needed to identify mechanisms at work, problems in place and eventually the dynamics of this new industry. The LTS framework offers a broader grasp of the infrastructural nature of AI as a technology, with more convenient categories to describe AI and measures to test empirically. We investigate how the implications of AI being an LTS entail the design of adequate public policies and firm strategies.

**Keywords:** artificial intelligence; large technical system; general purpose technology; infrastructural technology

**JEL Classification:** O33; L16; L52

We are grateful to Ed Steinmueller for his comments and the inspiring discussions shaping many of the arguments contained in this paper. We are also thankful to two anonymous reviewers of the SPRU Working Paper Series for their useful suggestions.

# 1 Introduction

Technological breakthroughs always acquire a life in their own right, as they beget novel perspectives or anticipate possible futures; doing so, they reshape our expectations and knowledge about potential states of the world. Given that, it is not surprising the increasing attention that social scientists have devoted to recent breakthroughs in Artificial Intelligence (AI). AI technologies are not an absolute novelty; rather, they experienced cyclical phases of hype and oblivion, often with scientific advances arriving with uneven time intervals in between, sometimes distant enough to create an impression of stagnation. The recent focus on AI grows out of the idea that 'this time is different' not only with respect to the previous phases of AI itself, but also in comparison with other technologies that are considered potential candidates to be at the core of technological revolutions. Regardless whether this time is really different, the idea that AI is a revolution is producing real–world effects. In fact, likely induced by the warmth of the current AI Spring that is reviving the curiosity and concerns frozen during AI Winters, many studies now focus on the technological features and socio–economic implications of AI.

The premise of human–level performance in many tasks creates expectations of AI's pervasive diffusion and, as a corollary, a key idea has been advanced: that AI is a so–called General Purpose Technology (GPT) (Goldfarb et al., 2019). Along with the literature on the Economics of AI, we think that AI should be assigned a special status among technologies. However, while acknowledging the transformative potential of AI, one can agree that tough every GPT is an influential technology, but not every influential technology is a GPT. So, does AI lie within or beyond the GPT definition? In the paper, we try to answer this question by mapping modern AI technology onto micro–characteristics and macro–effects assumed by the GPT framework. The result of this exercise suggests that despite AI having some touchpoints with a GPT, such as technological dynamism and innovation complementarities, equating AI and GPT is currently premature and, eventually, is likely to turn out as an incorrect definition of AI. The primary reason for this is that AI is qualitatively different from a stand–alone technology such as a GPT and instead resembles a system or infrastructural technology, approximating a *Large Technical System* (LTS) *in the making*.

LTS is one of the approaches to the study of technology that stresses the systemic — and thus structural and relational — properties of certain technologies. In fact, the perspective offered by LTS well suits the task of mapping system technologies that display wide reach.[1] The value of LTS theory for the Economics of AI lies in providing a 'thick description' of the technology: the vocabulary of notions it introduces is useful to unpack and understand the complementarities and tensions characterising infrastructural technologies and their development. This allows looking at AI from a novel angle offering a more insightful assessment of AI value, advancements and downsides, risks and benefits, room and tools for governance, opportunities and realistic trajectories of AI development.

Following the logic of the GPT part of the analysis, we introduce the LTS framework and map AI onto its building blocks, with a subsequent discussion of its goodness of fit. It should not be surprising that the LTS and GPT frameworks share or mirror some characteristics of

[1] An example of this approach at work for the case of telecommunication technologies is offered in Davies (1996).

each other, as both describe influential technologies, hence some similarity only speaks for the relevance of the comparison. As the differences in describing AI between the frameworks are of bigger interest, while evaluating AI as LTS we provide recollections of AI as a GPT for comparison and conclusion. As AI has just exited scientific laboratories and broke into the wild of commercial markets, it is yet in the making, and the mature form it will take is still to come. At such a key moment of AI development there are strong winds blowing in the direction of AI–infrastructure akin to the Internet, such as high returns on AI–based system–level substitution, concentrated market power among AI–producers, high costs of setting the system, etc. Neglecting these forces and treating AI in isolation from the rest of the system might lead to misplaced investments and dead–weight losses. Thus, the results of our analysis can be useful to researchers in the field of Economics of technological change and innovation as well as to policy makers, which might take–home from this study a better–suited, overarching framework to deal with AI.

The paper proceeds as follows. Section 2 places the first brick of the edifice by assessing whether it is correct to label AI as a GPT. Section 3 identifies which features of current AI map onto the LTS concepts. Section 4 derives implications for policy and strategy. Section 5 concludes.

## 2 Artificial Intelligence is a General Purpose Technology. Is it, really?

### 2.1 The 'next big thing': Artificial Intelligence

Scholars already consider AI the latest GPT. However, the search for a new GPT is not a novel endeavour. Over time, the title of general–purpose has been assigned to a non–negligible set of both narrow and broad technologies: electric dynamos (David and Wright, 2003), ICTs (Steinmueller, 2007; Strohmaier and Rainer, 2016), different computer platforms (Bresnahan and Yin, 2010), control technologies (Thoma, 2009), the Boyer–Cohen recombinant DNA technique (Feldman and Yoon, 2012), carbon nanotubes and nanotechnologies (Graham and Iacopetta, 2014; Kreuchauff and Teichert, 2014), and additive manufacturing techniques based on 3D printing (Choi, 2018). The rationale behind this race for GPT identification is one, common to all the studies: to find the 'next big thing' (Trajtenberg, 2018) that can induce profound socio–economic transformations while generating sizeable economic returns (Strohmaier et al., 2019). Going further back in the past, the search for the 'hopeful monsters' or macroinventions (Mokyr, 1990) that start new industries and pervasively transform the economy is a key feature of the literature on long waves, technological revolutions, and techno–economic paradigms (Perez, 2010). What these different approaches to technological change share is the idea that a certain technology might be the *primum movens* of long–term growth, development, and fluctuations (Kurz et al., 2018).

GPT and its diffusion is an ideal–type of a particular kind of technological change: one leading to the adoption of a radical innovation, used as a core component, across very diverse economic activities — in the limit, to ubiquitous adoption. Thus, to understand whether AI is a GPT is important, as that would provide insights into the future developments of the technology, the types of impact we can expect from it, and guidance for the design of policies to govern it.

**A short overview of AI**. As a first step of the analysis, we outline the framework through which we consider AI. This is necessary as AI is a 'suitcase word' (Mitchell, 2019) that densely packs an array of different meanings and interpretations. Mohamed et al. (2020) stress the dual nature of AI as *object* and *subject*: as object, AI is a set of technological artefacts; as subject, it is a 'portmanteau' of networks and institutions. Our analysis builds on this dual nature. In this section, we deal with AI as object and proceed through progressive approximations: from the philosophy of the technology to its particular instantiations. It is a useful exercise because the domain is dynamic and especially at the moment, when a handful of actors have entered the field with new products and new visions. In further sections, we move to AI as subject, building up the AI LTS from the core to its outskirts.

**Philosophy**. AI, being a technological mirror of ourselves, is inevitably compared to natural intelligence. The seemingly philosophical question of whether or not AI possesses a 'true' intelligence has very tangible technological implications in terms of, for example, engineering and programming. *Cognition and meaning understanding*, just to name two, are in fact the criteria and fields of ongoing research (see the new ICT taxonomy by Inaba and Squicciarini (2017)) that separate the so–called *weak AI* from *strong AI*. The distinction is based on the fact that the former only emulates intelligent behaviour, while the latter aims at re–creating it. While the emulation of intelligence is achieved using either rules of logic, heuristics, statistical learning techniques, or combinations of them, the question of how to re–create intelligence to reach true understanding by algorithms remains yet unanswered. Hence, the *current state of AI* belongs to the weak type. A relevant practical issue is that weak AI's reliance on statistical learning techniques entails risks for the deployment and usage of AITs. Incapable of general understanding, weak AI systems "have proven to be data hungry, shallow, brittle, and limited in their ability to generalize" (Marcus, 2020). Furthermore, neural architectures obtained through training can get obsolete, or can perpetuate biases that exist in the society (the 'garbage in, garbage out' principle). Such systems are vulnerable to (adversarial) attacks aimed at distorting or 'polluting' statistical (co–)occurrences in the data, teaching the system to behave oddly.[2] Moreover, several contingencies of the world with no clear (incomplete) ranking or dominance among alternatives remain challenging for weak AI to deal with (see for example the Moral Machine experiment (Awad et al., 2018)). Tweaking the algorithms in order to avoid these problems — correct for biases of data or society, create a decision–making routine for situations with no dominant strategy — and then retraining them entails high costs in terms of time, programming effort, computing power and energy, new tests, and environmental toll (Strubell et al., 2019)). In sum, current AI belongs to the weak AI domain and it has direct and tangible technical and societal implications with respect to which uses can be made of the technology and which risks it entails, and how to regulate the industry emerging around it.

**Approach**. We already mentioned statistical learning (including all: supervised, unsupervised or reinforcement) as a method to emulate intelligence. In general, there are two main approaches to AI: rule–based or symbolic approach (or good–old–fashioned AI, GOFAI) and

[2] A famous example illustrating such case is Microsoft's chatbot Tay: https://www.washingtonpost.com/news/the-intersect/wp/2016/03/24/the-internet-turned-tay-microsofts-fun-millennial-ai-bot-into-a-genocidal-maniac/.

statistical (data–driven) or connectionism.[3] In the symbolic approach an algorithm's search for a solution is driven by formal logic and explicit rules to deal with a given task, while connectionism uses statistical learning to infer implicit regularities from (un)structured data in order to perform a task. Currently, the latter is the prevailing approach to AI; it earned its fame due to the capability to learn from raw data without any predefined rules. This makes the connectionist approach autonomous, more flexible and effective in pattern recognition when compared to more rigid and bulky symbolic AI systems. However, connectionist AI algorithms, such as Artificial Neural Networks (ANNs), are tied to the task they perform (for instance visual recognition, language processing, or games — with only rare and partial exceptions, such as DeepMind's line of algorithms Alpha), so *they are function–dedicated virtual machines* (Boden, 2016).

**Technological constituents**. Regardless of the methodological approach, any AI technology necessarily consists of the following domains: (i) *algorithms or virtual machines*, (ii) *computing power* (and related physical devices delivering it), (iii) *data* and (iv) *domain structure*.[4] Domain structure indicates the problem environment and search space of actions an AI system is working with. Current AI technology applies an algorithm capable of learning to data in a given domain structure, requiring a certain amount of computation (which might vary substantially depending on the size of data and complexity of the algorithm and learning technique). In the context of this paper, listing these domains already makes clear that the set of actors involved in building AI systems should be extended from only algorithm–producing actors to a broader set including at least hardware producers, data providers, regulators, vendors, and buyers. Each domain is characterised by its own market structure, business models, regulations and pace of development. For instance, the semiconductor industry is a well–defined, consolidated manufacturing industry with a set of big players and high entry barriers, fuelled by capacity races and economies of scale (Steinmueller, 1992) and, from the technological side (until recently), by the Dennard scaling principle.[5] Differently, data is increasingly being considered as a commodity in recent years and debates on privacy and how to treat data are still ongoing (Savona, 2019). Despite the growing availability of data, compared to other domains, data markets remain opaque (Koutroumpis et al., 2020a).

**Functions and applications**. In order to complete the picture drawn in previous paragraphs, here we name some instantiation of current AI systems.[6] AI must not be conflated with either automation or robotisation, as AI plays only a specific role in each of these processes. Concerning automation, AI represents its high–end but is not a necessary condition for automation to take place; hence, AI is a subset of technologies associated with automation processes. With robotisation AI is rather in a parity relationship as robots, being mostly 'hardware', may or may not embody AI as a control method for physical work.

A last close up look at AI brings us to its practical functions and actual applications. Different functions like perception (visual and speech recognition), predictive analysis, communication

[3] Streams of research in AI are also partly moving towards hybrid approaches (Domingos and Lowd, 2019). We called our focus the 'current' understanding of AI as this is the one dominating research and societal attention at the moment.

[4] This description builds on Taddy (2019), but introduces some variations.

[5] The well–known Moore's law is a resulting expression of Dennard scaling as a physical principle.

[6] A good summary of AI techniques, functions and their applications fields is provided in WIPO Technology Trends Report 2019 on AI (WIPO, 2019).

(machine translation, information extraction) or control (robotics, facility–managing systems) might be involved separately or in combination in a variety of industries, from agriculture to advertising, Fintech, and even satellite communication (Ángel Vázquez et al., 2020). There are pioneer industries that deployed AI due to either the presence of specific functions that AI was capable of performing effectively or because the whole industry could come to existence thanks to AI; according to WIPO (2019), based on patent data, the top AI user–industries are transportation, telecommunication and life and medical sciences.

Taking stock, in this paper AI as object belongs to the *weak and connectionist AI*. In the next sections, we proceed by building around this core the economic, societal and institutional framework that constitutes AI as subject, and check this new systemic view of AI against the concepts of GPT and LTS.

### 2.2 Assessing the GPT nature of AI

**General Purpose Technologies.** Usually, to label a certain technology a GPT, characteristics of such technology are checked against the definitional criteria that describe what a GPT is. The most used definition is that of Bresnahan and Trajtenberg (1995), according to which a GPT is a technology that displays (i) *general applicability*, (ii) *technological dynamism*, and (iii) *innovational complementarities*.[7] General applicability captures the pervasiveness feature of GPTs. A GPT can be used as input or core component by a wide array of downstream industries or economic activities. Technological dynamism suggests that a GPT should display a steep learning curve in performance and/or efficiency — namely a fast ongoing improvement pulled by the enlarging downstream expenditure in the technology. Innovational complementarities, or innovation spawning, is instead the property that characterises GPTs from the perspective of innovative activities, rather than diffusion: GPTs are considered enabling mechanisms, as they induce or reinforce innovation incentives in the industries that use them as an input. Taken together, these characteristics illustrate a peculiar mechanism revolving around linked incentives: GPT producer and user sectors play a coordination game in which their optimal choices on technology are supermodular — the technological level of user sectors depends on GPT producers' product quality, and vice versa. A key feature highlighted by the GPT literature is that the impact of such supermodularity can have either a positive or a negative sign: due to the externalities coded into the GPT coordination game, both virtuous and vicious reinforcement cycles can occur, leading GPT development to feature multiple equilibria.

In what follows, we focus on the three core definitional characteristics of a GPT — general applicability, innovational complementarities, and technological dynamism, — together with two additional features, uniqueness and implementation lags, and assess whether they fit to describe the current AITs outlined in the previous subsection.

**General applicability of AI as a GPT**. The most characteristic feature of a GPT is its generality, or *pervasiveness*. A GPT becomes embodied in most of the technologies and used at

[7] The research programme on GPTs is characterised by a lack of coordination between different approaches, which has led to unresolved controversies (Bekar et al., 2018; Cantner and Vannuccini, 2012). For example, Lehrer et al. (2016) distinguish between nested 'mega GPTs' and 'anchor technologies', where the former are broad technological areas (such as ICT or nanotechnologies) and the latter are identifiable technologies (like semiconductor chips, or ERP software). From this descends the difficulty to draw a consensus definition of this family of technologies.

scale within the majority economic activities. This feature supports the claim that electrification or digitalisation are GPT–related processes: every device or product can be powered by forms of stored or non–stored electrical power, and most of the functions or activities conducted in an analog–mechanical manner (executed by relying on continuous input such as for example force or heat) can either be transitioned to digital (analog signals are replaced by discrete series of bits) or controlled digitally. However, the concept of pervasiveness carries a fundamental ambiguity, highlighted by Bekar et al. (2018) when they distinguish between *many uses* made of a particular technology, and technologies that are *widely used* across the economy. A technology with many uses is general purpose in nature, but that does not imply that it is also adopted at scale in the majority of economic activities; hence, the overall proportion of the economy that uses this technology might be small. In contrast, a single purpose technology can be an essential component in one or few industries. A GPT, in order to be pervasive, should permeate the economy in scale and scope — being widely used (at scale) in many uses (in scope).

It is undeniable that AITs are increasingly used in a disparate set of economic activities. What is remarkable of AITs is that they create *ex novo* activities in which they can be deployed — they kick–start new sectors and enable new products, e.g. autonomous vehicles. However, apart from some ex novo activity, one might argue that the majority of economic activities has only a limited reliance on AI. The fact that AITs' implementation at scale is localised in a few economic activities can be measured with respect to the following dimensions: penetration of (i) production processes, (ii) tasks within occupations, and (ii) overall adoption at the industry and firm level.

Looking at production processes, AI executes tasks that were already executed by capital, in particular ICT capital. The adoption of AI occurs through a replacement of existing software technologies with more sophisticated ones, those based on AI algorithms. This implies that AITs do not induce a substitution between production factors (capital for labor), and therefore the scale of task replacement is limited. Indeed, Bresnahan (2019a) suggests that AITs generate *system–level substitution*. System–level substitution occurs between production systems — for example online retail replaces brick–and–mortar one, automated user support or algorithmic fraud check replace the computer–aided but human–controlled version. Therefore, this process has to do with the introduction of new, more capital intensive 'production technology'; this includes the infrastructure underlying a firm's activities as well as its business model. In fact, AI–driven system–level substitution occurs in production processes that are already capital intensive pre–AI: these are a narrow set of economic activities or functions, oriented to consumer applications. Limited system–level substitution for AI contrasts the diffusion path at scale of the first wave of ICT (computers), and resembles the adoption dynamics of more recent ICT technologies, such as web and mobile applications: a targeted process of capital deepening in some activities (e.g. recommendation engines) leading to wide use by end–users and high returns but that leaves the rest of the economy substantially untouched.

Focusing on the occupation level, AI fails to permeate jobs at scale. Table 1 displays the share of AI–related jobs in the total jobs posted in the US for the year 2019, by NAICS 2–digit sector.[8] AI jobs are posted across a wide array of sectors (from information to construction), indicating

[8] The data is retrieved from the 2019 edition of the AI Index Report (Perrault et al., 2019) at the following link: shorturl.at/oAMU9

a spanning applicability of the workers' skills that are complementary to AITs. However, AI is not widely used: the share in the top–posting sector does not exceed 2.4% of total job posting, and is limited to values below 0.5% for half of the sectors considered. In line with this evidence, Acemoglu et al. (2020), find that while AI–related job postings accelerate, there is "no discernible impact of AI exposure on employment or wages at the occupation or industry level, implying that AI is currently substituting for humans in a subset of tasks but it is not yet having detectable aggregate labor market consequences". Exposure to AI affects some specific tasks within jobs, but not the occupational structure.

| **Industry** | **Share of AI jobs, %** |
|---|---|
| Information | 2.4 |
| Professional, Scientific, and Technical Services | 2.1 |
| Finance and Insurance | 1.3 |
| Administrative and Support and Waste Management and Remediation Services | 1.1 |
| Manufacturing | 1.1 |
| Management of Companies and Enterprises | 0.7 |
| Mining, Quarrying, and Oil and Gas Extraction | 0.6 |
| Agriculture, Forestry, Fishing and Hunting | 0.6 |
| Wholesale Trade | 0.5 |
| Educational Services | 0.5 |
| Public Administration | 0.5 |
| Retail Trade | 0.4 |
| Utilities | 0.4 |
| Health Care and Social Assistance | 0.2 |
| Real Estate and Rental and Leasing | 0.2 |
| Transportation and Warehousing | 0.2 |
| Other Services (except Public Administration) | 0.2 |
| Arts, Entertainment, and Recreation | 0.1 |
| Accommodation and Food Services | 0.1 |
| Construction | 0.1 |

*Source*: Perrault et al. (2019)

**Table 1:** Share of AI jobs posted (out of the total) by Industry, United States, 2019

For what concerns industrial connections, there are pieces of evidence that AI's diffusion among industries has a peculiar structure: despite being linked with many industries, these connections are shallow in the majority of cases. Using the expression of Bekar et al. (2018), AI has many uses, but is not widely used. For example, Prytkova (2021) considers the whole ICT system and estimates the scale and scope of industrial adoption of each distinct technology that constitutes the system, including AI. Figure 1 combines the results of Prytkova (2021)'s empirical analysis to illustrate industries' shallow reliance on AI. Figure 1a plots the change of scope (x–axis), i.e. number of AI's industrial linkages, versus the change in scale (y–axis), i.e. network centrality measure of AI as a technology connecting industries, between two periods — 1977–1990 and 2005–2020; the size of observations is the absolute value of the scale measure for the respective technology in the latest period. The reading of the figure indicates that AI acquired the largest number of industries between the two periods, but it is nowhere near to be adopted at scale. To reinforce the evidence, Figure 1b plots the average strength of industrial connections for each technology in the system; compared to other ICT technologies, AI ranks last.

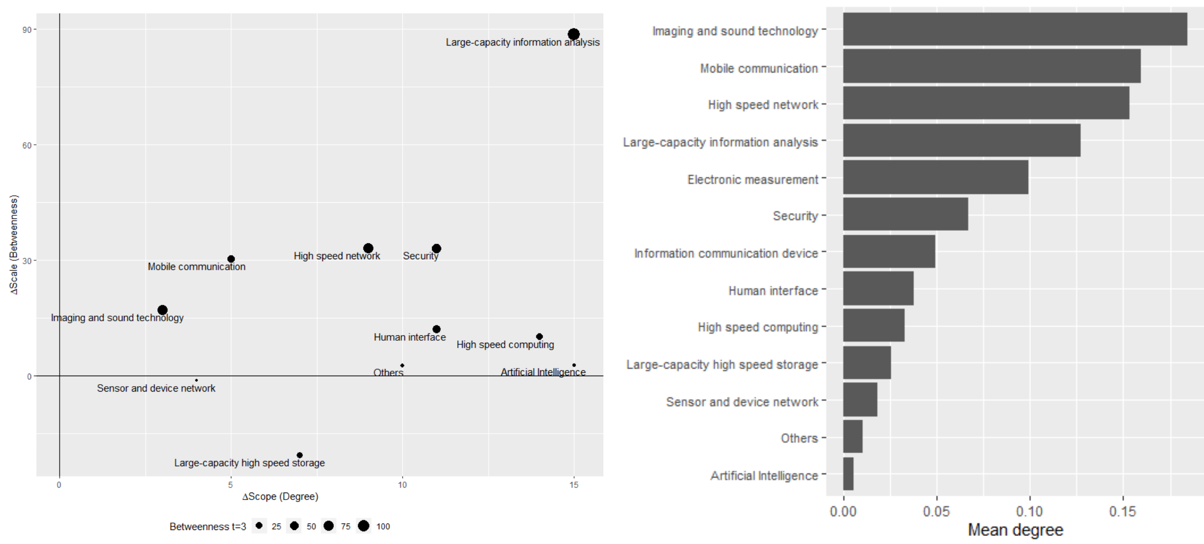


**(a)** Dynamics of AI's scope and scale

**(b)** Average strength of industrial connections

**Figure 1:** Industrial connections of AI (Prytkova, 2021)

In line with previous arguments and findings, Bresnahan (2019a) concludes that despite many uses, AI has gained a truly global reach mainly through one single economic activity: online retail and marketing (personalised advertisement and recommendation algorithms); even within this activity, diffusion is concentrated in and driven by a few, large and dominant actors — excluding them would further reduce the scale of AI adoption. AI applications that are not online retail or marketing sum up to a small fraction of the economy or to specialised and narrow user bases.

In terms of technology adoption at the firm level, Table 2 shows responses to the US Census Bureau Annual Business Survey 2018 — Digital Technology module, which provides information related to the year 2017 for all 3–digit NAICS sectors.[9] New business technologies use includes many AITs, such as machine learning, machine vision, natural language processing, and voice recognition (AITs are highlighted in bold). The table reports the share of surveyed firms that test, use or do not use these technologies; the evidence suggests that only a minority of firms report some use of AITs. In line with our argument, here we can interpret the many uses of AI as reported usage of several different AITs within firms, though not at scale — the share of firms testing or using AITs never goes above 3% of the total sample.

- In sum, across different dimensions of analysis — factors use and production technology, occupation structure and tasks, adoption at industry and firm level — AI achieved a limited penetration: a few economic activities rely on it at scale, while with the rest of the economy it developed a shallow connection. This means that AI is not pervasive in a GPT sense. However, can AI *become* pervasive as a GPT? Is the situation we outlined a temporary one, due to the infancy state of commercial AI? AI adoption is a moving frontier: theoretically, as more traditional operations are re–framed as prediction tasks, more activities could potentially be executed by AI. However, the way in which AI gains

[9]The data can be found here: https://www.census.gov/data/tables/2018/econ/abs/2018-abs-digital-technology-module.html.

| | (1) No use | (2) Testing but not using in production or service | (3) In use for less than 5% of production or service | (4) In use for between 5%-25% of production or service | (5) In use for more than 25% of production or service | (6) Don't know | Total share of use (including testing) (2)+(3)+(4)+(5) |
|---|---|---|---|---|---|---|---|
| Augmented reality | 80.0 | 0.3 | 0.3 | 0.2 | 0.2 | 19.0 | 1.0 |
| Automated Guided Vehicles or AGV Systems | 81.7 | 0.2 | 0.2 | 0.2 | 0.3 | 17.4 | 0.9 |
| Automated Storage and Retrieval Systems | 76.4 | 0.3 | 0.8 | 0.9 | 2.5 | 19.0 | 4.5 |
| **Machine Learning** | 79.3 | 0.5 | 0.8 | 0.7 | 0.8 | 17.8 | 2.8 |
| **Machine Vision Software** | 80.6 | 0.3 | 0.5 | 0.4 | 0.6 | 17.6 | 1.8 |
| **Natural Language Processing** | 81.1 | 0.3 | 0.4 | 0.3 | 0.4 | 17.5 | 1.4 |
| Radio-frequency Identification Inventory System | 81.8 | 0.3 | 0.3 | 0.2 | 0.3 | 17.1 | 1.1 |
| Robotics | 82.1 | 0.2 | 0.4 | 0.3 | 0.4 | 16.6 | 1.3 |
| Touchscreens/kiosks for Customer Interface | 77.8 | 0.7 | 1.3 | 1.2 | 2.3 | 16.6 | 5.5 |
| **Voice Recognition Software** | 80.8 | 0.6 | 1.0 | 0.6 | 0.5 | 16.6 | 2.7 |

*Source*: United States Census Bureau Annual Business Survey — Digital Technology Module 2018 (Table 3A: Business Technologies by 3–Digit NAICS for the United States and States)
*Note*: reference year 2017; numbers are totals for all sectors; number of firms surveyed: 4,618,795.

**Table 2:** Business Technology use in US firms (AITs highlighted)

scale is as an infrastructure, hence a measure like pervasiveness that is developed for stand–alone technological artefacts such as GPTs does not square well with AI producing little insights into the technology.

**Innovational complementarities of AI as a GPT**. Given its enabling nature, a GPT is expected to positively influence the *rate* of innovation in the GPT–user industries adopting it. The mechanism behind a GPT spawning innovation in downstream sectors is the so–called 'dual inducement' (Bresnahan and Trajtenberg, 1995). A dual inducement would occur when increasing the 'quality' of the GPT raises the curve of innovation returns for user industries; in turn, this raises the returns for the GPT sector to invest in GPT improvements. Dual inducement is typical of one–to–many architectures of technologies and industries resembling the broadcasting principle.

AI is certainly inducing higher rates of innovation: better AI algorithms are enabling more innovation in AI–using sectors, and the achieved positive results feedback on the incentives of AI–producing sectors to invest in further development of AITs. This description resembles a one–to–many (star) network, with pairwise connections between AI on one side and downstream sectors on the other side — as the stylised dual inducement suggests. In reality, for AITs the feedback is a systemic many–to–many process, with the whole collection of AI 'sibling' domains (hardware, software, data) connected to downstream sectors. AI evolves as a system, with innovation being 'pulled' by different downstream sectors; each sector calls for improvements in one or several AI domains that hinder its development. For instance, design of autonomous vehicles craves equally for more accurate algorithms because of their high stake loss function, faster processing and less energy consuming chips because of cars' battery capacitance, while more static applications like virtual assistants prioritise heterogeneity of computing and scalability. Even within the hardware domain, the established technological trajectory of semiconductors is being de–railed because of misaligned preferences among an increasing number of downstream sectors (Prytkova and Vannuccini, 2020). Another downstream sector of AI, the pharmaceutical and health industries, exert pressure on AI's development in two domains at the same time: algorithms and data. As for algorithms, the industry demands more explainable and at the same time better performing algorithms, that are usually associated with higher complexity and less explainability. As for data, the problem of availability of medical data to train and test algorithms' performance is tied to the debates on data privacy.

The role AI plays in innovation is broader than the one captured by GPTs' innovational complementarity. A GPT is a component that affects passively the innovation incentives of downstream sectors. Instead, AI actively participates in invention and innovation processes by creating information input: it can handle complexity ('needle–in–a–haystack' problems (Agrawal et al., 2018b)) and explore knowledge combinations in an automated manner, lowering search costs. While a GPT sets in motion a mechanism that raises the returns to innovation, AI directly helps innovating. From this perspective, AITs are invention machines (Koutroumpis et al., 2020b), and, thus, are closer to a so–called invention of a method of inventing (IMI; Griliches (1957)) than to a GPT. AI algorithms bruteforce the knowledge space (for example, corpora of annotated medical text) in order to identify potentially valuable associations and guide exploration. This has practical applications in business and in science. In business, AITs can intervene in product design and prototyping. In science, AI is increasingly used to aid the discovery of new drugs, materials, or biological structures such as the folding of proteins (Senior et al., 2020).

Despite the potential direct role in invention and innovation, AI is not displacing labour nor is used at scale even in this context. Bianchini et al. (2020) show that — at least for the Deep Learning technique and the case of health sciences — AITs do not yet work as a discovery 'autopilot' to explore and exploit the knowledge space. Rather, they remain an auxiliary research tool complementing existing scientific structures and practices.

- For AI, innovational complementarities have a networked, many–to–many nature: the inducement of innovation occurs among the (upstream) domains constituting AI as well as with (downstream) application sectors adopting AI. Moreover, AITs play a broader role than GPTs in inventive and innovative activities: rather than just influencing the rate of innovation, they are invention machines that actively participate in the process by automating the search for useful knowledge combinations and, thus, creating novel information input.

**Technological dynamism of AI as a GPT**. AI seems to display technological dynamism.

The performance of AITs compared to different benchmarks is improving quickly, so much to achieve above–human scores in some specific tasks (Eckersley et al., 2017). However, these tasks are yet the same that AI pioneers envisioned in the 1950s and 1960s, namely, tasks belonging to those fields in which human intelligence can be *emulated*, rather than generated *in silico*: pattern recognition, some perception (vision, speech), learning, bruteforcing of search through combinatorial spaces (Minsky, 1961). AI's technological dynamism is overestimated: the current improvement in performance is achieved through joint advances in different domains, as *a joint effort* taking place in the fields of computing, data and AI algorithms. This depends on the many–to–many structure of AI system, as we highlighted for the case of innovational complementarities. Thus, for AI, the feature of technological dynamism should be re–considered into what we call a *systemic technological dynamism*.

The point of systemic technological dynamism is that improvement is not isolated in one domain of AI, but involves and requires changes in many domains. Therefore, it is hard to trace where exactly the change started i.e. which domain 'decided' to improve. For illustrative purpose, we pick the algorithm domain of AI and unpack the system aspect of its dynamism. However, the reader must keep in mind that a networked system constantly circulates many changes from domain to domain, hence improvements occur as a result of purposeful action as well as 'fall from the sky'. In the algorithm domain, technological dynamism can be expressed through improvements in algorithm design, and is subject to a rather multifaceted dynamics. Algorithms are information goods; therefore, their production and consumption is shaped by incentives that are different compared to physical products (Shapiro et al., 1998). The kind of improvements of algorithm design are strongly driven by all sorts of efficiencies — computation per time, improvement in prediction per data batch, and trade–offs like bias–variance, complexity–explainability and other factors such as the ease of scaling up and costs of replication rather than by classic physical economies of learning (Nagy et al., 2013). Even at this fine–grained level of analysis, the constraints for improvement of algorithm design have a systemic nature: for example, if the possibility to improve the performance of algorithms depends on the programming environment used or (and) the hardware chosen, then algorithms evolution is function of strategic choices of the actors that control the programming environment and hardware production.

- In sum, technological dynamism in AI is not as rapid and linear as expected and occurs in bursts. The *systemic* nature of AI's technological dynamism manifests itself in the joint involvement of AI–related domains to achieve improvements.

**Uniqueness of AI as GPT**. Bekar et al. (2018) add another key criterion to the identification of a GPT. To be general–purpose, a technology should have "no close substitutes: (a) is unique — no other combination of technologies can produce an application; (b) without it the whole system (GPT and its application) would not work." The feature of uniqueness is helpful to discriminate between GPTs and what could be vaguely defined as radical innovations and important technological discontinuities. A GPT is identified as such because it does not have close alternatives: it is the sole and essential option to execute certain functions in user sectors. Uniqueness does not match well with current AI: in many economic activities, AI *is not unique*, but rather working autonomously, (can be) more efficient or/and precise at executing certain tasks compared to humans or to alternative techniques. For example, automated and algorithmic way of performing HR or business analytics can have a significant impact on firms' efficiency and economic returns, but this is not the only way to run these activities. Even though current AI has made its way into new applications and improved end–user experience, the whole system would not fall apart if AITs were to be rolled back. As illustrated earlier on, AI induces system–level substitution through capital deepening — in particular replacing older software technology with newer, AI–powered one, but the before–AI way of performing a task remains a close substitute, and in many cases yet a more reliable and precise one.

- In sum, AI has close substitutes for the functions it provides: it is not unique and rarely essential for the functioning of user sectors.

**Implementation lags of AI as GPT.** The diffusion of a GPT is expected to generate non–linear impacts on economic outcomes, in particular productivity (Jovanovic and Rousseau, 2005). GPTs do not necessarily produce these macroeconomic effects (Bekar et al., 2018). However, given their novelty and appeal to a variety of uses, it is possible that GPTs display implementation lags. The reason for it is that in order to exploit in full the pervasive potential of a GPT, resources already employed in productive uses need to be temporarily foregone and allocated to develop complementary assets (Brynjolfsson et al., 2021). For GPTs, implementation lags are demand–driven: in order to adopt it, GPT–users need to incur adjustment costs, among which those for organisational changes, capital investments, and development of skills to handle the new technology. In the case of AI, implementation lags are not necessarily driven by the same mechanism. The bottlenecks delaying AI implementation are mostly supply–driven: AI–producers need to obtain required inputs (data, hardware, and skills), set up production process and deliver a minimum viable product. For example, the collection of datasets for training AI models can take time and postpone the launch of AI products. AI producers can shorten the implementation lags by acquiring data on data marketplaces, exploiting cross–product data feedback loops, training their models using pre–trained models (teacher–student) or by "faking until they make it" using AI 'impersonators' (Tubaro et al., 2020) to buy time while training data is collected. These strategies are viable only in some cases and for some AI companies: data trade and access can be regulated; data feedback loops can be exploited almost exclusively by multi–product firms; the pre–trained models must be available, trustworthy and provide sufficient quality. Notwithstanding the potential remedies, and in contrast with the case of GPTs, bottlenecks for AI implementation remain a supply issue.

**Taking stock**. Is AI a GPT? Not exactly. AI is not pervasive in a GPT sense. It reaches adoption at scale only in a handful of industries, and even there diffusion is concentrated in and driven by a few large lead actors. Similarly to the Internet, AI provides an additional layer of functionalities to end–users rather than penetrating the economic structure: it is superimposed on existing systems. At the same time, its innovational inducement does not have a simple one–to–many nature and goes beyond a mere stimulation of the innovation rate in application sectors because AITs *participate* in and help navigating invention and innovation processes. AI technological dynamism is systemic and results from advances in interlocked sibling technological domains. Finally, AI has close substitutes, and its implementation can be subject to lags, but compared to a GPT the source of the lags for AI lie in the supply rather than demand side.

Using an econometric metaphor, *GPT is a misspecified model of AI*. The GPT misspecification originates from a potentially incorrect use of the included variables (functional misspecification) and, most importantly, due to omitted variables. The latter has two implications: first, it under– or overestimates of the importance of the included factors and, second, it misses a number of dimensions to represent AI adequately. Incorrectly specifying AI as GPT boils down AI to a poorly fitted, flat representation of what is instead a multidimensional complex phenomenon. Misspecifying an infrastructural technology as a single component will lead to incorrect inference and is likely to produce misleading predictions. It is possible to find a scheme that suits better the nature of AI. In the next section, we follow this route and try to look beyond AI–as–GPT.

# 3 Artificial Intelligence as a Large Technical System

## 3.1 Large Technical Systems

Large technical systems (LTS) are "spatially extended and functionally integrated socio–technical networks" (Mayntz and Hughes 1988). The notion belongs to the fields of sociology and history of technology, and science and technology studies. Compared to specific and isolated artefacts or technologies, LTS are 'system artefacts' or system technologies. Recognised examples of LTS are, among others, telecommunications, railways, energy supply and distribution systems. The prevalence of physical infrastructures among the mentioned examples of LTS does not exclude system technologies characterised by a higher degree of intangibility to be classified as LTS. In fact, Ewertsson and Ingelstam (2004) identify information–based LTS that contain both 'hard' and 'soft' components, such as radio and television distribution networks. Since the very introduction of the notion (Hughes, 1983; Hughes et al., 1987), the literature on LTS has investigated an array of issues characterising these system technologies, from definitional issues to the exploration of their dynamics and key actors. For the aim of this paper, the value added of the LTS theory lies in two dimensions: first, the outline of the different phases an LTS will experience from birth to maturity. Second, the identification of specific building blocks and driving forces that contributes to the formation and development of an LTS. These two dimensions are related, as different driving forces play a different role and have different relevance along the phases of LTS evolution.

The LTS phases originally singled out by Hughes et al. (1987) are (i) *invention*, (ii) *development*, (iii) *innovation*, (iv) *growth, competition and consolidation*, and (v) *technology transfer*. The latter is characteristic of LTS: technology transfer occurs when an LTS developed in a given context is replicated in other environments, and can happen in parallel to other phases. More recent work added new phases experienced by mature LTS, such a *stagnation, reconfiguration* and *decline* (Sovacool et al., 2018). Furthermore, Gökalp (1992) stresses how LTS develop by layering up over existing systems, creating a *superposition of systems* that shape an LTS configuration. The superposition of systems is characteristic of infrastructural projects and is an important feature to detect in an LTS. Complementary to the development in phases, a given LTS can be described as the result of a series of driving forces playing out to shape the infrastructural technology: *system builders, reverse salients, load factor, technological style*, and *momentum*.

**System builders**. System builders are the actors that strive to extend the reach of the system and perform the sociotechnical integration necessary to its deployment (van der Vleuten, 2009). These can be inventors–entrepreneurs or manager with engineering capabilities, individual actors or large firms. In different phases, system builders align the interests and objectives of the different actors involved, allowing an LTS to grow and achieve its goal(s).

**Reverse salients**. Reverse salients "are components in the system that have fallen behind or are out of phase with the others. Because it suggests uneven and complex change, this metaphor is more appropriate for systems than the rigid visual concept of a bottleneck. Reverse salients are comparable to other concepts used in describing those components in an expanding system in need of attention, such as drag, limits to potential, emergent friction, and systemic efficiency" (Hughes et al., 1987). Reverse salients, emerging from the uneven development of the system's components, are sources of critical problems and, given that problems are typically focusing devices (Rosenberg, 1969) to allocate innovative efforts, they are also potential loci of innovation.

**Load factor**. Load factor is "the ratio of average output to the maximum output during a specified period" (Hughes et al., 1987) and it is an indicator of performance, here meant as use or deployment of the technology at full potential over time. The distribution of load factor indicates when and where the system is under stress. Knowing that can guide investments in capacity expansions or adjustments, as well as policy interventions.

**Technological style**. As for the common use of the word, style indicates a type of fashion: the specific design of a particular LTS that descends from choices regarding which features are emphasised, and in which way. An LTS technological style emerges from the particular choice and combination of its elements, given their relative importance and the specific role they play in the whole system. LTS executing the same function and aiming at the same goal can differ in style in different contexts. For example, the organisation and control structure of energy distribution systems can change across countries while the fundamental function and goal they pursue are comparable.

**Momentum**. Momentum, or dynamic inertia, is the degree of autonomy the LTS acquires once it reaches a certain stage of development and a 'mass' in terms of relevance for the economic system. Systems with high momentum are less sensitive to pressures for change — they continue their 'motion' undisturbed.

The concept of system builder has mostly a social aspect, while reverse salient and load factor are dimensions of purely technological nature. Many of these concepts have closely related siblings in the field of economics of technological change. For example, reverse salients approximate bottlenecks; momentum approximates path dependence and cumulative change. However, their engineering or social flavour makes them more sophisticated categories to label complex phenomena, enriches the economics perspective and makes them useful to capture the features of system technologies that are uniquely embedded in specific epistemic communities, regulatory settings, and cultural contexts. A system builder can be an entrepreneurial actor, but also a carrier of a rare combination of technical and social skills (and, potentially, power). Momentum is close to path dependence, but path dependence is a process that emerges from chance and choices, while momentum is a later–phase property of a system that keeps existing and functioning due to 'mass' and acquired autonomy, thus refusing any role to chance.

The LTS categories are useful to guide the analysis of a given system technology. For example, one might want to know: where is the 'locus of control' in the system? Which actors store and hold the relevant technological (and market) knowledge to 'produce' the system technology? Who advances and builds the system out of its components? Who has power on the factors constraining the development of the LTS? Which elements of the systems and related actors can facilitate the process of convergence around standards and protocols in order to improve communication and control at large? What happens if the LTS becomes so large to be unmanageable? Joerges (1988) quotes Aristotle, reminding us that when things get too small or too large "they either wholly lose their nature, or are spoiled". A very timely point, when endless accounts of misuses, biases, discriminatory and malicious deployments suggest that we might be already spoiling AI.

### 3.2 Recognising features of LTS in AI

In Section 2.2, we checked some features of current AITs against GPT definitional criteria. The resulting picture suggests that AI substantially differs from a GPT, due to its rather infrastructural, distributed and heterogeneous nature. An alternative view on AI needs to encompass the whole circuit of actors and interconnections involved in its production and diffusion, their distinctive push and pull exert on the whole system, and a representation of how dispersed but linked activities influence the momentum of AI. We claim that LTS well–approximates the infrastructural nature of AI. To support this claim, we now identify element by element the LTS features in AI.

**AI is large**. LTS draws its specificity from the use of the attribute large. Following Joerges (1988) and Gökalp (1992), large can be considered in terms of territorial or user coverage, involving large–scale actors in the production of technology, or generating far–reaching socio–economic and/or environmental impacts. In this sense, large is used to label a technology that is encompassing, infrastructural, impactful, costly or global, or a combination of these properties. The attribute large is partially overlapping with GPT's pervasiveness, but it is broader, easier to measure, and fit for the purpose of describing an infrastructural technology. For example, a technology adopted only in one sector but at scale can be considered large. A different technology used marginally across a wide range of economic activities could be large as well, as overall it amounts to a wide reach.

The way in which large is measured in the LTS framework well–suits the measurement of the largeness of AI. Regardless of the few industries in which it originates, current AI spreads large in user base and territorial coverage given the widespread accessibility of its end–user applications. AI is large also because it is developed and promoted by a large community of actors (developers and vendors). The frontier of this large community is represented by large actors — large companies (the tech giants), national and supranational institutions, industrial consortia, global networks of Universities and organisations (and dedicated conferences). This creates a situation in which a large actor invests substantially into AITs and provides access to it to a large consumer base for sharing. For example, this is the case of sharing computing facilities and storage via cloud, or AI–powered software–based services (AI–as–a–service) such as visual recognition systems for airports. The economies of sharing (Shapiro et al., 1998) at work with AITs make the latter similar to classic LTS such as transport and energy supply systems. Finally, the societal traction of AI is large: "AI has seen itself elevated from an obscure domain of computer science into technological artefacts embedded within and scrutinised by governments, industry and civil society" (Mohamed et al., 2020); whole public opinions debate the changes AI will bring to contemporary societies, from its effects on employment, development and inclusiveness, its impact on minorities, and its environmental toll.

- In sum, AI is large according to various criteria identified by the LTS framework. This characteristic is better defined, inclusive and, hence, more convenient for both the identification of LTS and its empirical analysis.

**AI is a technical system**. AI is already implicitly considered a system from its very essential representation. The view of AI outlined in Section 2 helps to shed light on three constituent domains or subsystems that are key for the development of AI as LTS. First, the domain of AI algorithms that, in terms of actors involved and specific system builders engaged, is a subset of the software industry. Second, the domain of computation, in practice constituting a subset of the hardware industry. Third, the domain of data generation, collection, storage, analysis and transaction: data is collected and organised by public and private actors, globally and locally. As in a Venn diagram, at the intersection of these three domains one can find the state–of–the–art AI. These three domains are large in their own right according to the criteria we used early on: they are widespread (even if often invisible) in physical space, they contain numerous and large actors, and they are interwoven with and impacting socio–economic activities.

When discussing technological systems, Hughes et al. (1987) posits that they "contain *messy, complex, problem solving components*. They are both *socially constructed* and *society shaping*" (italics added). We unpack this statement to show how it tailor–fits to AI.

*The AI LTS is messy.* AI is still characterised by the turbulence typical of nascent industries, and uncertainties prevail with respect to its technological trajectories, its overall design, and its impacts. In the overall design of the LTS, one can devise alternative scenarios. As corner solutions, an AI LTS can be established either with a few large system builders dominating all the parts of the system or with an ecosystem of small actors scattered across domains. Intermediate settings, in turn depending on the direction taken by the regulation and governance of the LTS, can have large actors taking over some domains while leaving others untouched. Here, the relevant issue is to balance or align the societal and private interests of system builders and to identify important forking points in the path dependent process of AI development before the system gains so much momentum to become resilient to corrections. The very direction of evolution of AITs depends on the step–wise resolution of the current 'messiness'.

*The components of the AI LTS are complex and directed at problem solving.* Each of the domains of the AI LTS fits into the above statement. The case of AI chips production well captures the complexity of the hardware and computation domain of AI. Prytkova and Vannuccini (2020) summarise the trilateral frontier chipmakers address when developing their products: resolving a technical trade–off among delivering processing speed, energy efficiency, and heterogeneous computing. The data domain of the AI LTS is also complex: its current configuration is shaped by actors' competition to settle regimes of ownership and appropriation of data (Koutroumpis et al., 2020a). Spiekermann (2019) illustrates the structure of an *ideal–type* data marketplace that includes data buyers and sellers, the data marketplace (exchange) owner, and third–party service providers. AITs might be just tangent to the main goal of such data marketplaces (trading data), but perform an auxiliary function within this mechanism. From an AI–as–LTS perspective, the complexity in this domain arises from the fact that AI–producing and using companies can adopt different configurations: they can act as third parties only (AI–services providers), they can merge the role of third–party service provider and data buyer (e.g. using AI as a complementary technology to improve advertisement), and even layer–up the role of 'data exchange'. The latter is currently the case of Google, which owns a data exchange, uses AI to improve its products offer, and provides AI–based services (Srinivasan, 2019).

*The AI LTS is shaping society and it is socially constructed.* Harmful AI uses become increasingly evident the more AITs are implemented and turn into commercial and administrative tools. Concerns grow over specific applications of AI (e.g. face recognition), the ethics of algorithmic decision making, the safety of AI systems (e.g. to adversarial attacks), and the 'data colonialism' (Couldry and Mejias, 2019) premises on which these technologies are built, leading to a social pushback against harmful AI (Crawford et al., 2019). The acceptance or resistance to AI developments determines the social construction of this LTS (Mohamed et al., 2020). At the same time, the deployment of these technologies shapes society, in terms of perceptions (regarding, for example, the fears of AI–driven technological unemployment and widespread surveillance coexisting with the techno–optimism of grand opportunities on the brink of a fourth industrial revolution) and tangible implications. For example, companies started optimising their language and format in reports and disclosures, anticipating that these will be analysed by AI algorithms (Cao et al., 2020).

Finally, AI development clearly displays a *superposition of systems* (Gökalp, 1992). Current AI layers up on other LTS such as telecommunication and Internet infrastructure. Within public and private organisations, elements of their information systems are upgraded through AI–based capital deepening, while remaining integrated in existing organisational routines. For example, AITs integrate with existing database technologies, a trend already started decades ago (Brodie, 1989).

- In sum, computation (hardware), algorithms (software), and data domains are complex constituents of the AI system. The latter is yet messy, and hence contains the potential for different technological trajectories to unfold, leading to different designs of the system (e.g. distributed vs centralised, specialised vs general). This dynamic convergence to a more mature AI LTS happens through a simultaneous impact on and by the society.

Now that we have recognised the nature of LTS in AI, we can proceed with a more fine–grained analysis of the AI LTS by identifying the building blocks of LTS presented in subsection 3.1 in AI.

**AI system builders**. AI and its constituting domains are constructed by a variety of actors that actively initiate, support and shape developments of the system. The system builders in AI

LTS are *AI–producing and AI–using companies, dedicated regulatory bodies, industrial consortia, non–profit research organisations.* Every system builder exerts efforts to influence the selection of their priorities and problems for the system to implement and address. This can be done by trying to weave in a particular way the network of elements in the systems (an example are tech giants hiring AI pioneers and leading figures to lead their AI programmes) or by forcing the very system to converge on new standards, protocols and shared practices. The latter can be achieved by making obsolete or ineffective the status quo through, for instance, forking decisions (Simcoe and Watson, 2019).

Consider AI–producing and AI–using companies. The latter are mostly recipients of the novel technology and are experimenting with ways to integrate AI in their 'runtime' (Iansiti and Lakhani, 2020), while the former are proper system builders. AI–producing companies have the power to design the system and to decide which bridges between actors and subdomains to build or cut–off. AI–producing companies are an ecosystem of firms that conduct AI research, develop AI solutions, and participate in the AI value chain (Tubaro et al., 2020).[10] Among them, key system builders are the already mentioned tech giants, and established software and hardware companies[11], vendors, startups and platforms[12]. In terms of main line of business activity and industrial classification (NAICS), the majority falls under the codes 'software publishers' (49%) and 'Computer Systems Design and Related Services' (17%).[13]

One example that illustrates how system builders exert their power as such is the ongoing issue revolving around handling harmful AI. In this context, current commercial system builders have supported the establishment of ethical boards and voluntary guidelines rather than regulation. While regulation would impose common and accountable rules on the development of the system, commitment–based solutions can be considered forms of strategic concessions[14] to other AI LTS stakeholders (in particular regulators, consumers of AI services, and the society at large). In practice, these initiatives represent a 'seductive diversion' that allow AI–producing companies to show engagement while retaining full power over the design of the system. Another leverage AI system builders acquire with their role is their 'knowledge holding' (Steinmueller, 2006). In the production of AITs, novel know–how is created and knowledge about it settles in the hands of AI–producers. Through this process, AI system builders become knowledge gatekeepers that can facilitate the diffusion of knowledge and expertise through co–invention activities as well as strategically withhold it. For example, machine learning platforms — open or closed source (Isdahl and Gundersen, 2019) — might improve access to and reproducibility of AI solutions, however at the cost of product development knowledge being held for a larger share by the platform owner.

System builders in AI are changing over time, and their variety is increasing. AI companies superseded individual AI pioneers and universities' computer science departments; at the same time, they are accompanied by governments and 'lateral' organisations. The latter are,

[10] For an attempt at identifying AI–related companies, see https://cset.georgetown.edu/research/identifying-ai-related-companies/

[11] These companies take the lion's share in AI patenting. According to Van Roy et al. (2020) top AI patenting firms (in the period 2010–2016) have their main activity in 'Manufacturing of electronic equipment' and 'Information and communication'.

[12] https://www.venturescanner.com/category/artificial-intelligence/

[13] https://cset.georgetown.edu/wp-content/uploads/CSET-Privately-Held-AI-Companies-by-Sector.pdf

[14] https://onezero.medium.com/the-seductive-diversion-of-solving-bias-in-artificial-intelligence-890df5e5ef53

for example, advocating to make the system more inclusive and less harmful (e.g. AI Now), pursuing technical advancements through non–profit organisations (e.g. Open AI), facilitating coordination on principles and standards (e.g. the Partnership on AI), or stressing the importance of getting prepared to the emergence of strong AI (e.g. the Future of Humanity Institute). Another type of system builders, currently less empowered than the ones mentioned earlier on, are the (platform) workers that support the deployment of AI systems and that are subjects of processes of 'heteromation' (Tubaro et al., 2020).[15] These workers operate at the margins of AI and fill gaps in the working of the technology — they run the so–called 'AI last mile', either fuelling the data necessary for the training of algorithms, verifying their performance or even emulating the results of AI systems.

**AI reverse salients**. As the system scales and becomes larger, tensions appear. These fault lines are the reverse salients of the system. One recurring source of reverse salients in AI are the system's scarce resources in AI's domains: *being a nascent industry, AI lacks input resources from its domains. The shortage is relative among domains, i.e. the worst performing domain is a source of reverse salient, which can be of quantitative or qualitative kind: delivering an insufficient amount of an input resource, or a qualitatively unfit input.* This holds back or derail the evolution of AI LTS. Quantitatively speaking, AI is data–hungry but other resources whose demand grows faster than supply can become constraining factors as well: among them, AI labour and AI–programming skills, and management trained to lead AI–powered companies. Qualitatively speaking, computation is a reverse salient that falls behind because the dominant design of chips in the semiconductor industry doesn't match the way modern AI operates, even though some alternative trajectories emerge (Hooker, 2020; Prytkova and Vannuccini, 2020). Atop of the purely technical challenge, there is also a techno–economic one: the competition between cloud and edge modes of organisation of the computing infrastructure. The cloud mode entails allocation of resources with priority on coverage and speed of access networks, and on computing capabilities of cloud–providers. The edge mode emphasises connectedness (which is not the same as coverage), compatibility and computing capabilities of edge devices. Evidently, the choice of the prevalent mode defines the chips' design that will get the lion's share of R&D efforts. The resolution of this qualitative reverse salient from the computation domain will shape the AI LTS with regard to the organisation of computation: inside the chip among its components as well as inside the industry among producers and users.

Another example of qualitative reverse salient is the lack of a well–designed and regulated architecture for data troves, originating from scarcities in the data supply. AI systems can rely (i) on public dataverses and open data or (ii) on the supply of data from data marketplaces. These two alternatives carry with them two philosophically competing views: in the first, the reverse salient is addressed through the creation of common assets; this can give a more democratic style to the AI LTS, but can also open room for nationalistic interpretations of the idea of data 'sovereignty'. The second view might lead to an 'oligopolistic' style of the AI LTS, with

[15] Introduced by Ekbia and Nardi (2017), the concept of heteromation documents the shift, in some sectors, from technologies of automation, in which machines take over tasks from humans, to technologies of heteromation, in which tasks at the margin of value creation that are devoted to the management and update of automated systems are externalised to humans. In this sense, rather the producing unemployment or the end of work, the introduction of automated systems actually increases labour demand; this is however concentrated in labour intensive micro-work activities, with consequent detrimental impacts.

a few powerful players shaping the playground at their own advantage. Compared to the 'AI commons' scenario, the oligopolistic one might hasten the growth and impact AITs, but can lead to a more unequal distribution of returns.

Reverse salients emerge also in the domain of AI algorithms. One lies in the proliferation of AI software and programming environments, slowing down the convergence towards a dominant design. Part of the community of AI developers urges technical improvements through recognised contests dedicated to different AI problems[16], open–source platforms to assist the coherence of the community and the development of cross–compatibilities, the establishment of standardised libraries and programming frameworks, and more fundamental theoretical and technological advances (Ben-David et al., 2019; Geirhos et al., 2020; Marcus, 2020).

Another reserve salient is overspecialisation among AI algorithms. Despite AI algorithms become increasingly capable (see Hernandez and Brown (2020) for an assessment of algorithmic performance and efficiency trends), the tendency for ad hoc solutions remains. The reason for that lies in the pursuit of a sole criterion of performance (or its derivatives), namely, out–of–sample accuracy of prediction. The development of algorithms proceeds along this criterion and hence relies heavily on the intensive margin, a trend succinctly expressed as "the bigger the better" — whether bigger refers to the size of a model, of data or of computing power. Figure 2 supports this statement plotting accuracy versus model size for two different AI tasks, visual recognition and natural language inference.

The upper panels of Figure 2 show decreasing returns to number of parameters in both tasks: as the number of parameters in a model grows, the corresponding gain in accuracy is getting smaller. Black lines represent borders of the Pareto–Koopmans criterion (PKC) (Bogetoft and Otto, 2010); at the intersection of the PKC borders lies a model with the highest accuracy to size ratio, i.e. productivity. The empty second quadrant indicates absence of more efficient observations; after the intersection point the returns on model size are decreasing in terms of accuracy of prediction. The lower panels of Figure 2 show that the linearised relation between model size and productivity is strong for both tasks. A deviation upward of the fitted line would indicate a higher return on the number of parameters than expected for a corresponding model size, but there are no such deviations. These results illustrate the claim of algorithms' development along the intensive margin, with accuracy improving at slowing down rate at the expense of accelerating model size. For example, in 2020 GPT–3 model by OpenAI has 175 billion parameters, 100 times larger than models launched two years before. The new frontier of model size, achieved in 2021, is Google's Switch Transformer, featuring more than 1.5 trillion parameters and hence jumping in size by a factor of 8.6 in one year. Techniques like parameter pruning, quantisation, transfer learning, and the usage of lower precision arithmetic might be steps towards more efficient models.

Reverse salients originated in the domain of algorithms have implications for the hardware domain: ad hoc AI algorithms appeal to smaller demand and have short–lived returns, quickly becoming obsolete. At the same time, the design and production of a chip that caters the needs of an ad hoc AI solution has high sunk costs. Therefore, the resolution of reverse salients in the algorithm and hardware domains is entangled, and both remain in a turbulent state until

[16] Numerous examples can be found here: https://paperswithcode.com/datasets and, more hardware–oriented, the MLPerf training benchmark: https://mlperf.org/.

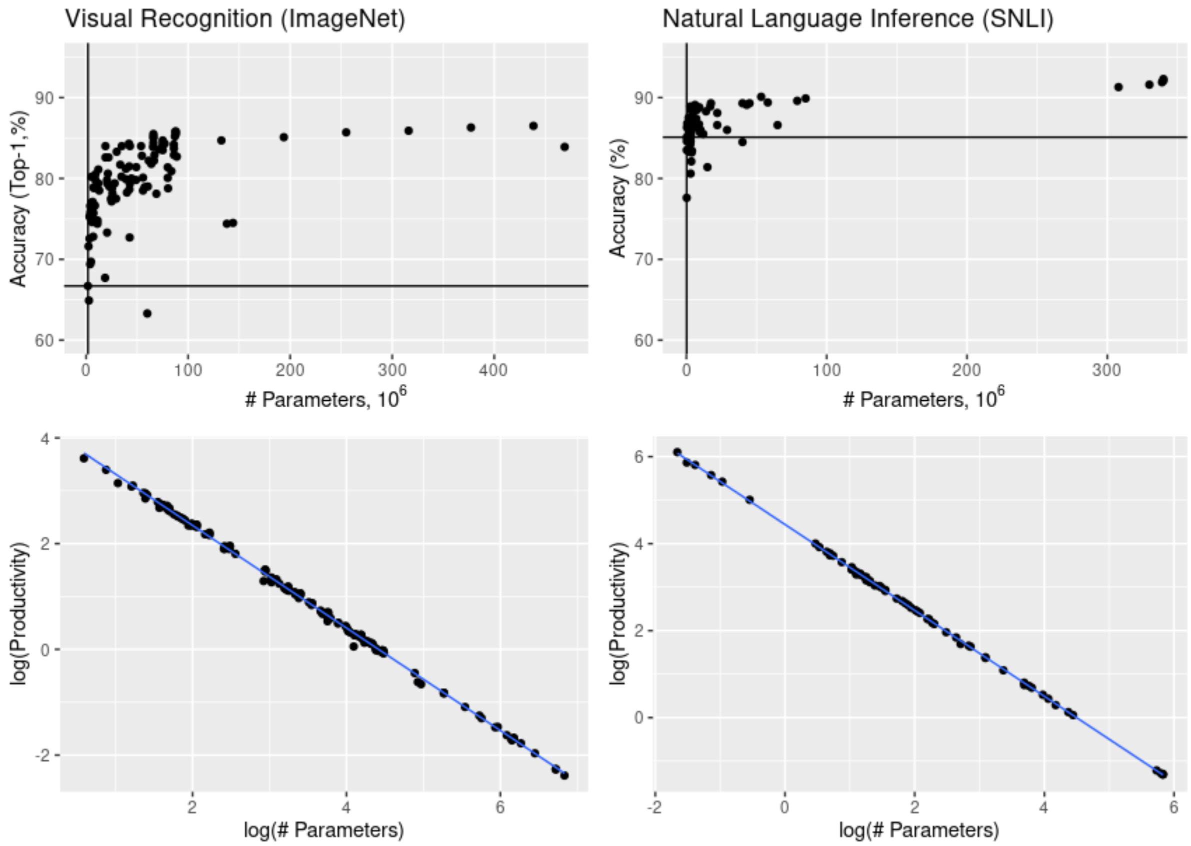


**Figure 2:** Models' retarding performance and decreasing efficiency in visual recognition and natural language inference tasks

Note: Each observation is a model with reported specifications and performance trained on the same data sourced from the respective benchmark test. *Upper panels*: accuracy of out–of–sample prediction plotted versus number of parameters in a model measured in millions. An observation at the intersection of the Pareto–Koopmans criterion borders (black lines) marks the model with the highest return on number of parameters. *Bottom panels*: linearised relation between model size and productivity measured as ratio of accuracy to number of parameters labelled as productivity. *Left panels*: visual recognition task using data of ImageNet competition. *Right panels*: natural language inference task using the SNLI corpus.

a dominant design emerges in either of the domains. Though connected, the development of the two domains remains driven by rather separate criteria, treating each other as a source of constraint for its own performance rather than a tandem partner with a common goal. The current AI shock shed light on this issue, exposing the shortcomings of such divided approach; there appeared several studies that call for an expansion of performance criteria into various *joint efficiencies*, beyond separately accuracy in the software domain and processing speed in the hardware domain (Chen et al., 2019; Hooker, 2020; Prytkova and Vannuccini, 2020; Strubell et al., 2019). In line with our discussion on systemic technological dynamism, lack of awareness about the joint nature of improvements brought separate technological trajectories to their exhaustion and further opportunities are recognised in joint efforts of the software and hardware domains.

**AI momentum and load factor**. A system can be improved by solving tensions created by reverse salients, producing what Sahal (1985) defines *learning by scaling*. The learning occurs by accumulating understanding on which changes (innovations) to implement, and how to redistribute the load in the system while it scales up. In other words, knowing the load factor guides directed interventions to address salients and to allow for improvements *ceteris paribus* the level of technical performance. On top of that, Nightingale et al. (2003) suggest the notion of 'economies of system' to explain the gains a system can enjoy by redistributing activities according to the load factor, dynamically balancing the stress.[17] Economies of system in the AI LTS would occur by rearranging the structural dependencies among its elements when some of them develop unevenly or are overloaded. For example, the shift to federated learning architectures (Li et al., 2020) would represent a system re–arrangement towards a design potentially capable of addressing the computation–related reverse salient: this would be done by distributing workload over the networked components rather than leaving few giant actors to route (and control) finite computing power in the cloud.

Finally, the growing number of actors jumping on the bandwagon of AI successes, the grandiose media coverage of AI advances (in particular for what concerns language models and generative models of speech — the so–called conversational agents like Project Debater of IBM (Slonim et al., 2021)) and the expectations of further ubiquitous diffusion of AI build up a strong momentum for the AI LTS. However, expectations can work in both positive and negative direction. On the one hand, they channel large investments in AI R&D by public and private system builders. On the other hand, the expectations of a large and ubiquitous impact of AI risk remaining unfulfilled: sustained commercialisation and growing competition among system builders make them race against each other, undertaking myopic steps in AI development leading to short–term payoff. Stagnating diversity of AI research is among the early signs of such dynamics (Klinger et al., 2020). As the expectations that a new AI winter might be at the horizon start to be considered plausible and the AI hype slows down, the momentum of the system might follow a similar path.

### 3.3 AI LTS: State–of–the–art

Having identified the technological and non–technological features of LTS in AI, we can proceed with a description of state–of–the–art AI LTS: the phase of development the system has achieved, the boundaries that confine the system, the mechanisms of control currently in practice, the distinctive style emerging, and finally the goal or a main function the system embeds.

**Current Phase**. The 'invention' phase of AI is a contested territory, as the very understanding of what AI is shifts over time; this is why in Section 2 we offered a view of current AI. We can claim that following the impressive results in the ImageNet visual recognition competition in 2012 and the subsequent media interest in AI — mostly due to the shadows AI seemed to cast on the future of work — the AI LTS went through the phases of invention and development. The current state of the AI LTS is now in–between the phases of innovation and growth, competition and consolidation, with commercialisation accelerating its pace and increasing technology transfer from academia to business, including a sizeable talent drain of professors and graduate students (Zhang et al., 2021). The very process of growth by expanding to novel application fields generates continuous feedback into the phases of innovation and development. Technology transfer has also accelerated with the increasing efforts of national, supranational and sub–national institutions to govern AI developments as the technology has

[17] The notion shares similarity with that of architectural innovation (Henderson and Clark, 1990), where improvements in performance are achieved by changing the arrangement of the constituents of a technology while keeping its function invariant.

acquired geo–strategic significance.[18] While in the process of innovation and growth, an interesting question regards whether a process of technological convergence is taking place within the AI LTS. Technological convergence, a concept introduced by Rosenberg (1963), is a form of 'upstreaming', a process occurring when an activity embedded within diverse sectors or/and tasks exhibits some common features and principles that eventually maturates and unbundle into a fully–fledged sector on its own.[19] We see signs of this process at work in the evolution of the AI LTS: AI–producing companies — key system builders — emerge as specialised suppliers of AI–as–a–service tools, business automation services (e.g. recruitment), scientific discovery tasks ('science–as–a–service') and data analytics tasks. The success of a few system builders to impose their template on the working of the AI sector will influence the evolution of the whole LTS infrastructure in the future, set the stage for novel reverse salients to appear, and orient policy priorities.

**Boundaries of the AI LTS**. While the AI LTS evolves, it passes from one phase to another and its boundaries change: the system grows larger, usually in an uneven manner. This creates the problem of where the AI LTS 'begins and ends'.

The span of the AI system was at first limited to a pure scientific territory at the intersection of computer science and psychology; then, the entire landscape of business actors that could commercialise AI has been incorporated into the system.[20] The first generation of AI–based products and services has been introduced, and their overlooked flaws — especially with respect to AI safety and ethics — became a fertile soil for the next cohort of actors that address these flaws to enter the system. Progressively, in the described manner, more areas were incorporated in the AI LTS advancing its frontier. The application fields of medical and biological sciences, robotics, automated decision making in business and public administrations and the military sector have recently stepped into the LTS and add new forces to the complex casting of the direction and pace of AI evolution.

**Control over AI LTS**. A relevant issue related to the shifting boundaries of the AI LTS is the possibility that diseconomies of scale exists, leading the system to lose internal coherence and to face 'crises of control' while growing (Beniger, 1986). Under control we understand mechanisms of coordination on progression and management of the deployed applications. The issue of control arises when these mechanisms are non–existent or function inefficiently. Nightingale et al. (2003) study and provide examples of innovations in control technologies as devices to retain control of LTS as they scale up. The issue of control, however, seems to be less relevant for the AI LTS: given the heterogeneity of its components, the system is less coherent as an ensemble compared to more homogeneous LTS. This could be the result of AI being still a 'young' LTS, with a yet fluid distribution of power over the system. In a sense, among LTS, AI is closer to the Internet than to integrated transport systems. As with the Internet, AI displays a mix of

[18]The reader can refer to the OECD AI Policy Observatory (https://oecd.ai/) or to Nesta AI Governance database (https://www.nesta.org.uk/data-visualisation-and-interactive/ai-governance-database/) for a collection of AI–related policy initiatives and institutional strategies.

[19]The separation of chip production from the rest of the computer industry or the classic formation of the machine tool industry described by Rosenberg (1963) can be considered cases of technological convergence. When technological convergence takes place, a new industry can emerge upstream, producing generic technologies that suit a wide variety of downstream purposes.

[20]Earlier AITs, such as expert systems, extended already in the 1980s the boundaries to Industry and commercialisation activities (Nilsson, 2009).

centralised and decentralised mechanisms of control and a layering up of commercial and non–commercial areas of development (Greenstein, 2020). Coordination among the system builders is achieved through the convergence on standards and interfaces, which is a non–frictionless process. From the perspective of control, in AI, there are not yet agreed standards for progression of the system, nor an essential need for the centralised maintenance of coherence among the system's parts; also, the actors do not have to be explicitly aware of the infrastructural nature of the AI LTS in order to conduct their operations. Therefore, the minimum requirements to make the AI system working and to avoid the system falling apart are lower than for other infrastructural technologies, while the social impact is potentially larger for AI. This does not mean that AI evolution will continue to follow the same loosely coordinated path: there are coalitions of system builders in formation, supporting the idea of development and implementation of standards (both technical and ethical) to 'govern' AI. Such coalitions have different goals and different shapes, according to the kind of system builders they aggregate: workers and end–users of AI (Crawford et al., 2019), strategic alliances, intergovernmental initiatives (such as the 'Partnership on AI', and fully–fledged consortia. This activism represents the search for directions of progression of the system. In the current phase, the system already gained significant momentum, but 'diseconomies of scale' have not occurred and internal coherence has not been upset due to the compartmental structure of AI LTS. Instead of the 'death' of the AI system, the failure of control mechanisms could steer AI towards towards detrimental directions at crucial points of its evolution path, where detrimental is intended here in the sense of harmful or exploitative of a share of its actors or users.

**Technological style**. The AI LTS will display different technological styles in different environments. The specific design of the system constituents and the strength of their interdependence vary according to the priorities, strategies, and decisions taken by system builders, regulators, and users. This becomes evident when considering national (and supra–national) implementations of AI systems. The 'division of labour' and the direction of AI development and deployment depend in part on the structure of AI– and data value chains (Tubaro et al., 2020), but it can be strongly affected by government strategies, resulting in rather distinctive styles. The first distinction can be between the technological style that shapes around the effort of *prolonging the 'age of discovery'* and the one focused on *accelerating the switch to the 'age of implementation'* (Lee, 2018). The first style stresses the role of continuous innovation in AI techniques. The second considers the main AI breakthroughs as already achieved and aims at scaling them up by accelerating diffusion and experimenting with applications to capitalise on them.

AI styles can form under influence of diverse system builders. Let's consider the role of the state in giving the AI LTS its style. Technological style can emerge as a result of the particular policy levers and priorities the regulator decides to pursue. This is reflected in public budget allocations that can channel funds to AI through the university system, the military sector[21], or directly to private actors — for example in form of financial support to AI startups. Beyond the sheer amount of expenditures and the broad direction imposed on the system's evolution, different technological styles for AI LTS can emerge as a result of the specific tools of technology

[21]See the 2021 Report of the US Department of Defense or National Security Commission on Artificial Intelligencehttps://www.nscai.gov/2021-final-report/

policy used (Steinmueller, 2010). Here, *top–down command and control policy* actions share the stage with more *bottom–up governance initiatives*. Horizontal interventions, such as the design of regulatory frameworks against AI harms, misuses and biases, are part of a specific style. The creation of new dedicated institutions (as it has been debated regarding the possibility to create a federal robotics commission in the US — see Calo (2014)) and intermediate bodies to facilitate coordination in the system is another, potentially complementary, option. Examples of policies that can influence technological style are the efforts made by governments to attract, retain and develop AI talent through the visa regime,[22] or the alignment of macro policy levers (e.g. immigration and trade policy) with AI–related strategic priorities. A relevant case of the latter option are export controls policies targeting the semiconductor industry, as this is the producer of key components for AI–tailored hardware and its productive capabilities are a fundamental strategic asset. The use of policy levers in strategic technologies such as AI is not a novelty: the just cited semiconductor industry has been subject of trade policy interventions to shield domestic companies against emerging competitors (Langlois and Steinmueller, 2000). This point highlights the tight link existing between the actions leading to the development of a technological style and the competition between institutional actors at the international level. AI becomes ones of the territories over which geopolitical forces play, so much that some authors discuss whether an AI 'arms race' is ongoing (Asaro, 2019).

**Goal orientation of the AI LTS**. As a final step in our analysis, we assess whether AI has an identifiable goal orientation. LTS are goal–oriented systems; this means that all components coordinate — easily or at the cost of frictions, negotiations and forced adjustments — to achieve an overarching aim. This characteristic is easy to detect in LTS such as transport systems or water distribution networks (the goals being, respectively, mobility and water supply), while it is less evident for telecommunications, the Internet or, indeed, AI. The reason why the goal is less evident for the last three cited systems is because not only they are large, but also highly heterogeneous and less prone to have centralised control being exerted on the LTS participants. We posit that AI LTS is also goal–oriented, even though the goal might not be clear yet to all the actors involved. We suggest that the goal of the AI LTS is a 'cybernetic' one: *to re–domain the 'fabric of control' of socio–technical systems based on human decisions into an automated one*, starting with the transformation of existing activities into instances of adaptive prediction.

Within the closed system of steps required to complete a given task (Agrawal et al., 2018a), currently, AI is focused on perfecting pattern recognition and prediction capabilities. In general, prediction serves the needs of the decision–making stage of task performance, as "each prediction task is a perfect complement to a decision task" (Agrawal et al., 2019). Essentially, any prediction shrinks the search space that will inform action and lead to a desirable outcome. Up to now, AI is allowed, trusted or able to take decisions only in a few contexts. Simulated, virtual or physically–confined environments such as games, trading and test sites (e.g. for cars, drones and robots) are the contexts at the forefront of AI performing as an autonomous decision maker. In all the other contexts, judgement or decision–making remain overwhelmingly in the hands of the human counterpart. However, the consideration of a task as a closed system suggests looking at the entirety of the process of task performance. To establish full cybernetic control over the

[22] https://cset.georgetown.edu/research/immigration-policy-and-the-global-competition-for-ai-talent/

performance of a task, AI has to permeate each stage necessary to that task's execution: (i) elaboration of input data (pattern recognition, prediction), (ii) judgement and decision–making, (iii) action and feedback. An example of task controlled by AI along all stages is the industrial control system of cooling facilities in Google's data centres that went completely autonomous in 2018.[23] In general, it is clear that the achievement of the overarching goal of cybernetic control requires the maturity of multiple technologies and institutions, and their coordination. To accelerate or steer this process, reverse salients (technologies, mechanisms, institutions) that are falling behind and holding back AI can be identified adopting the view of AI being an infrastructural technology already now and an LTS in the future. In sum, to use the terminology of Flueckiger (1995), the goal of the AI LTS is to shift further the balance from economies based on operations of transformation to economies based on operations of control — and to automate these.

# 4 Implications for Policy and Strategy

Seeing AI as an LTS rather than a GPT has important implications for policy and strategic decision–making. The core argument here is that the rationale for and the essence of intervention differs between the AI–as–GPT and AI–as-LTS case. To illustrate that, we can compare how the focus of policy might change by changing the categorisation of AI. When a technology is identified as a GPT, the rationale for intervention lies in market failure. The key issue is the under–production of the GPT technologies due to the distributed nature of downstream innovative efforts, which would require coordination. Fixing a coordination failure in the GPT case means kick–starting the dual inducement mechanism, raising the rate of investments in innovation until to foster positive feedback. In this context, public procurement and contract spending can emulate, substitute or subsidise downstream demand. When a technology is an LTS, coordination issues extend beyond simple incentive formation, and become a matter of joint design and production of the whole network of technologies involved in the system. From this perspective, failures take the form of system or orchestration failures, with actors failing to develop the necessary ties and alliances to strike a balanced development of the system (Robinson and Mazzucato, 2019; Schot and Steinmueller, 2018). Rather than facing a stagnating innovation rate, reverse salients appear locally and slow down or disable the whole system, making it work inefficiently or even miss its goal(s) entirely. In system technologies, the source of failure might be located within one component, distributed among several components or even be the very disconnectedness of the system itself. For an LTS, the correct identification of reverse salients and the detection of their composition and reach across the system is a primary step to undertake. Once diagnosed, the task becomes to devise a strategy to tackle the problematic areas of the LTS network, inducing desirable effects and preventing the side effects of the 'treatment'.

From this perspective, the AI LTS requires policy makers to get to know the specificity of the system under consideration: who are the system builders, where are the boundaries of the system, which mode of control is at work at a given moment and locality, how the load factor is measured and distributed. Policy makers must adopt systemic thinking to acquire awareness

[23] https://www.technologyreview.com/2018/08/17/140987/google-just-gave-control-over-data-center-cooling-to-an-ai

of the state of the LTS, its current phase and potential paths of evolution, in order to inhibit detrimental or catalyse dormant useful activities, components and actors, fill gaps and missing links in the system, rebalance control or redistribute load factor, and in general to decide if to opt for command–and–control types of intervention or to prefer indirect forms of governance. Depending on which reverse salient is addressed, policy can opt for a different recipe of science, technology, industrial and competition policy tools (Steinmueller, 2010).

To show how strategy and policy can be discussed from the AI–as–LTS perspective in details, we take the AI reverse salient related to data and summarise dimensions relevant to AI deployment and upon which policy makers can act. Over the last 10 years, we observe a growth of business models that are reliant on the monetisation of data. The diffusion of the Internet and the globalisation of markets at the same time made possible an unprecedented expansion of the consumer base, a boom in the amount of offers from businesses of all kinds, and drastically lowered the related (information) search costs and the cost of tracking the consumption behaviour (content, goods, services, etc.) of online users (Goldfarb and Tucker, 2019). Atop of this abundance of data, new market opportunities for businesses that collect, store, structure and elaborate the data rapidly grew: online databases, search engines, consulting firms, digital platforms, software management systems and many other examples of data–fuelled business models. This is a key transformation: where there is data, there will be AI. AI has the potential to spread into applications where data (i) is generated and can be collected in sufficient amounts, and (ii) its structuring and elaboration creates value–added for the business. These conditions shape the data reverse salient and expose the non–pervasive character of current AI.

**Getting the data**. First, in order to deploy AI to support any given application, an established and systematic process of data collection is required. In other words, the implementation of AI requires a meaningful representation of business processes (essential or not for a firm) in data — namely, their digitisation. This is why pioneering industries in AI adoption are the likes of Fintech and logistics, which are characterised by highly digitised and measurable processes and had forms of algorithmic automation and optimisation already in place. The so–called 'Deep Learning revolution' stands precisely in the fact that it provided an effective tool to process raw unstructured data e.g. images, video, audio, making this activity cheaper (and thus economically viable) and less labour– and time consuming. Doing that, Deep Learning expanded the set of tasks that can be solved by AI algorithms. Deep Learning made possible to exploit troves of raw data that were already out there, waiting for an algorithm to harness them. An example is AI–based visual recognition, which emerged as a novel function applied to medical imaging records for diagnostics in many medical disciplines.

The existence of data does not automatically make the case for an AI application. Sometimes data might exist but its accessibility could be either hindered, inefficient or even welfare–damaging. This is partially due to unresolved data ownership and absence of mechanisms such as data markets to coordinate data supply and demand which would ensure the lawful and effective exchange of data ownership rights. An insightful summary of the situation with data markets is expressed in a quote of Edward Snowden: "there is no property less protected and yet no property more private than data" (Snowden, 2019). In some applications, data is a mere representation of an environment's state or processes (e.g. temperature control in data centres).

However, when data is an imprint of activities conducted by actors, individuals or organisations that are external to owners of AITs, then data might be considered as a property of the actors that created it (Jones and Tonetti, 2020). Said differently, when data is a public good, ownership issues do not emerge, while the elaboration of data, which has the nature of a private good, requires solutions that address simultaneously consensual data transfer and privacy concerns (personal data that owners might either sell at a very high price or not to sell at all).

- In sum, the collection of data that reflects business processes including demand's feedback loops and establishment of data markets is a necessary though not sufficient prerequisite for AI deployment.

**Monetising the data**. Second, to persist being used as a useful technology within an economic activity, data elaboration performed by AI has to bring returns. The value of data elaboration can lie in harnessing otherwise unmanageable amounts and complexity of data or (and) detecting patterns that humans cannot identify. Retrieving information about, for example, highly non–linear relations between a set of covariates and whether or not a person has clicked on an ad is undoubtedly a useful insight, but in order to systematically turn this information into a profit a firm has to build a sustainable business model to monetize on it. Monetisation strategies can vary across applications, which in turn are characterised by different payoffs from the implementation of AITs. For example, for online retail, the monetisation strategy would involve the structuring of pricing and versioning of the offer given the association revealed by data elaboration. This strategy allows obtaining profit directly and from each offer independently. Differently, an AI algorithm that controls an industrial robot through the processing of sensory data and producing an adequate response in order to perform a routinized task creates value added that is more implicit and grows in a non–linear way with the scale of deployment of the technology.

- In sum, all kinds of data elaboration done by AI has to produce either valuable/unique intermediate result in the firm's production process or contribute to a valuable offer to the consumers, in both B2B and B2C markets, to ensure retention and generate profit.

**Investing in assets**. Third, sustaining the monetisation strategy requires investments into complementary assets of some kind. The costs of primary collection or acquisition of data from third parties (e.g. the purchase of database licences, cookies or data appends — see Bergemann and Bonatti (2019)), the storage within a firm or purchasing cloud space in order to further elaborate the data with AI, and even contracting micro–work to conduct data annotation (Tubaro et al., 2020), constitute yet another part of the data–related reverse salient. Thus, depending on the revenue from AI–based activity, a firm has to choose between investments into the development of AI systems at least in part in–house (including all domains — data, hardware and software) or into partnership with AI–provider along the AI value chain. The choice between the two alternatives is intertwined with the control aspect of AI and it shapes the distribution of market power among system builders in the nascent AI industry. Obviously, small and medium–sized enterprises tilt toward outsourcing option to minimise costs. Moreover, even big companies for which AI performs not a core but a side function would be prone to purchase customised but ready–made AI solution in a package, benefiting from sharing the risks and legal responsibilities with the developer. Indeed, among AI–users the emergent strategy of 'join–and–share' AI–as–a–service solutions due to the high costs of every component of AI systems steers AI development towards a form of infrastructure, with the most powerful system builders (AI–producers) meticulously building and gathering pieces of the infrastructure together. The burden of high costs is coupled with cross–domain network effects. For example, depending on the application, the nature of data might vary — pixel matrix for images, text corpus for legal disputes, or panel data for consumer databases. This affects the choices and developments in the hardware domain (bandwidth capacity, memory size and placement, parallel or sequential processing and so on), programming framework (programming language, libraries) and algorithms themselves (loss function, optimization procedure). Together, the initial costs of implementation and cross–domain network effects increase switching costs of an alternative to any component and lead quickly to hard lock–ins for both supply and demand in the software and hardware domains. The result of this dynamics is a trend of over–specialisation in both domains, as we discussed in Section 3.2. Investments in more versatile and heterogeneous hardware and algorithms is a long–term strategy, but it has a longer period before returns start and is associated with uncertainty regarding adoption, making such innovation trajectories affordable only to a minority of (rather large) system builders.

- In sum, AI adopters make a choice on how to deploy AI–based solutions and invest in the respective complementary assets. This creates a demand–pull effect steering the innovative efforts of AI–producers further along existing technological trajectories. The opportunity costs in this situation might be substantial, as alternative trajectories are locked out by prohibitively high switching costs.

Eventually, two incentives reinforce each other: competition among AI–producers makes them sensitive to demand's (AI–users) needs, while demand is following the visibility of commercial value to sustain its strategies to monetise on adopted AI solutions rather than the technical superiority of these solutions in the long term.[24] The task for policy–makers is to make sure that arguments related to high costs and strong network effects are not used as a justification to tilt the development of the infrastructure towards an inefficient realisation. In the case of AI, inefficient in technological sense might mean avoiding following the already–mentioned principle 'the bigger the better' in terms of ever–increasing size of data, amount of compute, complexity of algorithms, number of processors and so on for the sake of marginal improvement in performance (refer again to Figure 2, for instance). The fundamentally statistical nature of current AI will always strive for more data as a safe solution not only to achieve better representativeness of a given sample, but also to train deeper ANNs following the trend to increase the size of algorithms. In socio–economic sense, an inefficient instantiation would drain resources and resemble a skewed representation of stakeholders' interests — AI–users, AI–producers, society, public institutions — creating dead–weight losses, violating rights, damaging competition, and producing an asymmetric distribution of gains. R&D investment in scalable AI techniques like federated learning and neural network compression, and hardware technologies such as platform

[24]This dynamics has been observed for Machine tools industry by Rosenberg (1963) and for ICTs by Bresnahan (2019b).

chips and edge computing can soften hard lock–ins and create ways out through compatibility with already existing components of the infrastructure.

Given the discussion above, we can outline a set of insights for policy–making: in order to cultivate technological opportunities to implement AI, policy attention can be directed to address the grey areas of data creation, collection and distribution. A way to do that is to assess how it has been done within the pioneer applications of AI. In particular, focusing on firms, this entails filling gaps such as developing the capabilities to digitise a firm's processes, organising their systemic and structured execution, and creating a digital twin of a firm's activity to be analysed with AITs. From the firms' perspective, the business models that monetise on AITs must be flexible to avoid being locked in solutions offered by dominant actors in monopolistic or oligopolistic markets. From the policy perspective, attention should focus on monitoring, detecting and regulating the whole *network* of AI–related markets, to ensure the conditions for fair competition among system builders, and to lower the cost of exploration and support of alternative technological solutions and partnerships. This would nurture an ecosystem of actors and technologies contributing to the transition to a more distributed mode of control over the AI LTS.

Overall, if AI is an LTS then policy design should be inspired by the priorities set by the LTS framework. Examples of these priorities are: (i) the balanced construction of the system, for example by supporting the development of AI talent, identifying and suggesting new components for the system based on relatedness, providing resources and facilities for experimentation; (ii) curbing the monopolisation of resources in the hand of a few actors across the fundamental domains of AI ensuring equal access for all system builders; (iii) pushing for inclusive or public models of governance by pursuing the identification of technical and non–technical standards.

# 5 Conclusion

Artificial Intelligence is generally considered a breakthrough that is technologically revolutionary and, often, also philosophically existential; it is capable of reshaping societies and economies while at the same time offering a mirror to look inside ourselves and our human intelligence. Adapting an expression of Norbert Wiener, if AI is ready to 'usurp specifically human functions', then our priority should be to understand which is the 'purpose put into the machine' to avoid unintended consequences that can result from delegating tasks to AI.

Indeed, AI has the potential to influence many real world processes. But the very nature of current AI is less romantic than what is usually depicted, even though its impact can still be transformative. Much of the hype around AI is a case of what Braitenberg (1986) called 'the law of uphill analysis and downhill intervention', according to which humans tend to overestimate the complexity of a mechanism, guessing its internal structure from observation. The human–level performance displayed by AI on certain tasks can indeed hide the 'Clever Hans' nature of these technologies. In fact, current AITs are essentially a new wave of ICT technologies. At the moment, they are usually brittle and function–specific algorithms that are used as 'prediction machines' because they are effective in identifying associations and extracting patterns.

In this paper, we focused on the essential mechanisms of AI and offered a novel perspective on its nature — a key exercise, as AITs are increasingly embedded into products and service, commercialised and used in a wide range of applications. In particular, we tested in details the consensus idea that AI is a general purpose technology by evaluating how GPT definitional characteristics fit the features of AI. Our conclusion is that it is premature to consider AI a GPT. This is not because AI is a technology just emerging, and thus *not yet* a GPT, but instead because the GPT 'suit' is structurally inappropriate — and namely too flat — to dress AI. AI is not a stand–alone technology as GPTs are, but a system technology that displays infrastructural properties: it has a dual nature, as a technological artefact and at the same as a socio–technical network.

AI shares some features with GPTs (for example innovational complementarities and technological dynamism), but these have a qualitatively different nature in the AI case. The very differences of AI from the GPT benchmark are what carries useful information. For example, we establish the stylised fact that, at different levels of analysis, AI is not pervasive in a GPT sense: it has many uses, but it is not widely used in the majority of economic activities — it is not as ubiquitous as computers are. Even in the few industries in which it is adopted, diffusion is concentrated in and driven by a few large lead actors. Similarly to the Internet, AI provides an additional layer of functionalities to end–users, and for this reason it spread large in user base and territorial coverage when embedded in final products and services; besides that, AITs leave the rest of the economy almost untouched.

Being a GPT requires penetration in scope and scale, and AI might never reach that status, remaining confined in activities where it is useful the most. Transformation occurs at a deeper level: the level of systems, where AI is implemented as an additional layer in a system to become an autonomous and capable (and, hence, perfect) tool of cybernetic control. In this way AI becomes an infrastructural technology: superimposed on existing technological layers, and, through that, interwoven into the economic structure.

If GPT is a misspecified model of AI, is there a better model around capable to capture the nature of an infrastructure technology? Our bet is on the concept of Large Technical System and the framework it offers to describe system technologies. In the analysis, we mapped the LTS building blocks on AI. This allowed us to identify AI's system builders, reverse salients, momentum and load factor, as well as to derive useful insights on AI phases of development, boundaries, control, technological style, and goal orientation. Furthermore, we applied the AI–as–LTS scheme to the issue of policy and strategy and showed that the rationales for intervention and the types of policy actions differ substantially between the AI–as–GPT and AI–as–LTS interpretation. As an example, we focused in details on the reverse salient related to data to illustrate the many tensions emerging within this domain of AI.

With our study, we contribute to the nascent Economics of AI and, more generally, to that part of the Economics of technological change and innovation interested in uncovering the structure of technological breakthroughs. Using the LTS framework, we extend the reach of economic analysis of AI to the neighbouring fields of sociology of technology and science and technology studies. As a result, we were able to offer a novel understanding of infrastructural and system technologies through the case study of AI.

As a new fully–fledged industry is rapidly forming around AI, a correct mapping of the technology and its complex nature is necessary to avoid misunderstanding its trajectory, misallocating resources dedicated to its progress, and harmful developments. Understanding AI means understanding its fundamental fabric and design principles: how a system technology is engineered by different actors in a dynamic 'workspace', which forces shape its path of development, and how these same forces can be steered in a direction that contributes to the common good.

## References


Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). Ai and jobs: Evidence from online vacancies. Technical report, National Bureau of Economic Research.

Agrawal, A., Gans, J., and Goldfarb, A. (2018a). *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Agrawal, A., McHale, J., and Oettl, A. (2018b). Finding needles in haystacks: Artificial intelligence and recombinant growth. Technical report, National Bureau of Economic Research.

Asaro, P. (2019). What is an artificial intelligence arms race anyway. *ISJLP*, 15:45.

Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J.-F., and Rahwan, I. (2018). The moral machine experiment. *Nature*, 563(7729):59–64.

Bekar, C., Carlaw, K., and Lipsey, R. (2018). General purpose technologies in theory, application and controversy: a review. *Journal of Evolutionary Economics*, 28(5):1005–1033.

Ben-David, S., Hrubeš, P., Moran, S., Shpilka, A., and Yehudayoff, A. (2019). Learnability can be undecidable. *Nature Machine Intelligence*, 1(1):44–48.

Beniger, J. R. (1986). The control revolution: technological and economic origins of the information society.

Bergemann, D. and Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11:85–107.

Bianchini, S., Müller, M., and Pelletier, P. (2020). Deep learning in science. *arXiv preprint arXiv:2009.01575*.

Boden, M. A. (2016). *AI: Its nature and future*. Oxford University Press.

Bogetoft, P. and Otto, L. (2010). *Benchmarking with DEA, SFA, and R*, volume 157. Springer Science & Business Media.

Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. MIT press.

Bresnahan, T. (2019a). Artificial intelligence technologies and aggregate growth prospects.

Bresnahan, T. and Yin, P.-L. (2010). Reallocating innovative resources around growth bottlenecks. *Industrial and Corporate Change*, 19(5):1589–1627.

Bresnahan, T. F. (2019b). Technological change in ict in light of ideas first learned about the machine tool industry. *Industrial and Corporate Change*, 28(2):331–349.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of econometrics*, 65(1):83–108.

Brodie, M. L. (1989). Future intelligent information systems: Ai and database technologies working together. In *Readings in artificial intelligence and databases*, pages 623–641. Elsevier.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–72.

Calo, R. (2014). The case for a federal robotics commission. *Brookings Institution. Brookings.*

Cantner, U. and Vannuccini, S. (2012). A new view of general purpose technologies. Technical report, Jena Economic Research Papers.

Cao, S., Jiang, W., Yang, B., and Zhang, A. L. (2020). How to talk when a machine is listening: Corporate disclosure in the age of ai. Technical report, National Bureau of Economic Research.

Chen, Y.-H., Yang, T.-J., Emer, J., and Sze, V. (2019). Eyeriss v2: A flexible accelerator for emerging deep neural networks on mobile devices. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems.*

Choi, J. (2018). The rise of 3d printing and the role of user firms in the us: evidence from patent data. *Technology Analysis & Strategic Management*, 30(10):1195–1209.

Couldry, N. and Mejias, U. A. (2019). Data colonialism: Rethinking big data's relation to the contemporary subject. *Television & New Media*, 20(4):336–349.

Crawford, K., Dobbe, R., Dryer, T., Fried, G., Green, B., Kaziunas, E., Kak, A., Mathur, V., McElroy, E., Sánchez, A. N., et al. (2019). Ai now 2019 report. *New York, NY: AI Now Institute.*

David, P. A. and Wright, G. (2003). General purpose technologies and surges in productivity. *The economic future in historical perspective.*

Davies, A. (1996). Innovation in large technical systems: the case of telecommunications. *Industrial and Corporate Change*, 5(4):1143–1180.

Domingos, P. and Lowd, D. (2019). Unifying logical and statistical ai with markov logic. *Communications of the ACM*, 62(7):74–83.

Eckersley, P., Nasser, Y., et al. (2017). Eff ai progress measurement project.

Ekbia, H. R. and Nardi, B. A. (2017). *Heteromation, and other stories of computing and capitalism.* MIT Press.

Ewertsson, L. and Ingelstam, L. (2004). Large technical systems: A multidisciplinary research tradition. In *Systems Approaches and Their Application*, pages 291–309. Springer.

Feldman, M. P. and Yoon, J. W. (2012). An empirical test for general purpose technology: an examination of the cohen–boyer rdna technology. *Industrial and Corporate Change*, 21(2):249–275.

Flueckiger, G. E. (1995). *Control, Information, and Technological Change*. Number 6. Springer Science & Business Media.

Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., and Wichmann, F. A. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673.

Gökalp, I. (1992). On the analysis of large technical systems. *Science, Technology, and Human Values*, pages 57–78.

Goldfarb, A., Gans, J., and Agrawal, A. (2019). *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press.

Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.

Graham, S. J. and Iacopetta, M. (2014). Nanotechnology and the emergence of a general purpose technology. *Annals of Economics and Statistics/Annales D'Économie et de Statistique*, (115/116):25–55.

Greenstein, S. (2020). The basic economics of internet infrastructure. *Journal of Economic Perspectives*, 34(2):192–214.

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica, Journal of the Econometric Society*, pages 501–522.

Henderson, R. M. and Clark, K. B. (1990). Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative science quarterly*, pages 9–30.

Hernandez, D. and Brown, T. B. (2020). Measuring the algorithmic efficiency of neural networks. *arXiv preprint arXiv:2005.04305*.

Hooker, S. (2020). The hardware lottery. *arXiv preprint arXiv:2009.06489*.

Hughes, T. P. (1983). Networks of power: Electric supply systems in the us, england and germany, 1880-1930. *Baltimore: Johns Hopkins University*.

Hughes, T. P. et al. (1987). The evolution of large technological systems. *The social construction of technological systems: New directions in the sociology and history of technology*, 82.

Iansiti, M. and Lakhani, K. R. (2020). *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*. Harvard Business Press.

Inaba, T. and Squicciarini, M. (2017). Ict: A new taxonomy based on the international patent classification.

Isdahl, R. and Gundersen, O. E. (2019). Out-of-the-box reproducibility: A survey of machine learning platforms. In *2019 15th international conference on eScience (eScience)*, pages 86–95. IEEE.

Joerges, B. (1988). *Large technical systems: the concept and the issues.* Wiss.-zentrum für Sozialforschung.

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–58.

Jovanovic, B. and Rousseau, P. L. (2005). General purpose technologies. In *Handbook of economic growth*, volume 1, pages 1181–1224. Elsevier.

Klinger, J., Mateos-Garcia, J., and Stathoulopoulos, K. (2020). A narrowing of ai research? *arXiv preprint arXiv:2009.10385.*

Koutroumpis, P., Leiponen, A., and Thomas, L. (2020a). Markets for data. *Industrial and Corporate Change*, 29(3).

Koutroumpis, P., Leiponen, A., and Thomas, L. D. (2020b). Digital instruments as invention machines. *Communications of the ACM.*

Kreuchauff, F. and Teichert, N. (2014). Nanotechnology as general purpose technology. Technical report, KIT Working Paper Series in Economics.

Kurz, H. D., Schütz, M., Strohmaier, R., and Zilian, S. (2018). Riding a new wave of innovations. *Wirtschaft und Gesellschaft-WuG*, 44(4):545–583.

Langlois, R. N. and Steinmueller, W. E. (2000). Strategy and circumstance: The response of american firms to japanese competition in semiconductors, 1980–1995. *Strategic Management Journal*, 21(10-11):1163–1173.

Lee, K.-F. (2018). *AI superpowers: China, Silicon Valley, and the new world order.* Houghton Mifflin Harcourt.

Lehrer, M., Banerjee, P. M., and Wang, I. K. (2016). The improvement trajectory of pcr dna replication and erp software as general purpose technologies: an exploratory study of ‘anchor technologies’. *Technology Analysis & Strategic Management*, 28(3):290–304.

Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60.

Marcus, G. (2020). The next decade in ai: four steps towards robust artificial intelligence. *arXiv preprint arXiv:2002.06177.*

Minsky, M. (1961). Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30.

Mitchell, M. (2019). *Artificial intelligence: A guide for thinking humans.* Penguin UK.

Mohamed, S., Png, M.-T., and Isaac, W. (2020). Decolonial ai: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, pages 1–26.

Mokyr, J. (1990). Punctuated equilibria and technological progress. *The American Economic Review*, 80(2):350–354.

Nagy, B., Farmer, J. D., Bui, Q. M., and Trancik, J. E. (2013). Statistical basis for predicting technological progress. *PloS one*, 8(2):e52669.

Nightingale, P., Brady, T., Davies, A., and Hall, J. (2003). Capacity utilization revisited: software, control and the growth of large technical systems. *Industrial and Corporate Change*, 12(3):477–517.

Nilsson, N. J. (2009). *The quest for artificial intelligence.* Cambridge University Press.

Perez, C. (2010). Technological revolutions and techno-economic paradigms. *Cambridge journal of economics*, 34(1):185–202.

Perrault, R., Shoham, Y., Brynjolfsson, E., Clark, J., Etchemendy, J., Grosz, B., Lyons, T., Manyika, J., Mishra, S., and Niebles, J. C. (2019). The ai index 2019 annual report. *AI Index Steering Committee, Human-Centered AI Institute, Stanford University, Stanford, CA.*

Prytkova, E. (2021). Ict's wide web: a system-level analysis of ict's industrial diffusion with algorithmic links. *Available at SSRN.*

Prytkova, E. and Vannuccini, S. (2020). On the basis of brain: Neural-network-inspired change in general purpose chips. *SWPS.*

Robinson, D. K. and Mazzucato, M. (2019). The evolution of mission-oriented policies: Exploring changing market creating policies in the us and european space sector. *Research Policy*, 48(4):936–948.

Rosenberg, N. (1963). Technological change in the machine tool industry, 1840-1910. *Journal of economic history*, pages 414–443.

Rosenberg, N. (1969). The direction of technological change: inducement mechanisms and focusing devices. *Economic development and cultural change*, 18(1, Part 1):1–24.

Sahal, D. (1985). Technological guideposts and innovation avenues. *Research policy*, 14(2):61–82.

Savona, M. (2019). The value of data: Towards a framework to redistribute it. *SWPS.*

Schot, J. and Steinmueller, W. E. (2018). Three frames for innovation policy: R&d, systems of innovation and transformative change. *Research Policy*, 47(9):1554–1567.

Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., Qin, C., Žídek, A., Nelson, A. W., Bridgland, A., et al. (2020). Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710.

Shapiro, C., Carl, S., Varian, H. R., et al. (1998). *Information rules: a strategic guide to the network economy.* Harvard Business Press.

Simcoe, T. and Watson, J. (2019). Forking, fragmentation, and splintering. *Strategy Science*, 4(4):283–297.

Slonim, N., Bilu, Y., and Alzate, C. (2021). An autonomous debating system. *Nature*, 379–384.

Snowden, E. (2019). *Permanent record*. Macmillan.

Sovacool, B. K., Lovell, K., and Ting, M. B. (2018). Reconfiguration, contestation, and decline: conceptualizing mature large technical systems. *Science, Technology, & Human Values*, 43(6):1066–1097.

Spiekermann, M. (2019). Data marketplaces: Trends and monetisation of data goods. *Intereconomics*, 54(4):208–216.

Srinivasan, D. (2019). Why google dominates advertising markets. *SSRN*.

Steinmueller, W. E. (1992). The economics of flexible integrated circuit manufacturing technology. *Review of Industrial Organization*, 7(3-4):327–349.

Steinmueller, W. E. (2006). Learning in the knowledge-based economy: the future as viewed from the past. *New Frontiers in the Economics of Innovation and New Technology. Essays in Honour of Paul A. David. Cheltenham, UK: Edward Elgar*, pages 207–238.

Steinmueller, W. E. (2007). The economics of icts: Building blocks and implications. In *The Oxford handbook of information and communication technologies*.

Steinmueller, W. E. (2010). Economics of technology policy. In *Handbook of the Economics of Innovation*, volume 2, pages 1181–1218. Elsevier.

Strohmaier, R. and Rainer, A. (2016). Studying general purpose technologies in a multi-sector framework: The case of ict in denmark. *Structural Change and Economic Dynamics*, 36:34–49.

Strohmaier, R., Schuetz, M., and Vannuccini, S. (2019). A systemic perspective on socioeconomic transformation in the digital age. *Journal of Industrial and Business Economics*, 46(3):361–378.

Strubell, E., Ganesh, A., and McCallum, A. (2019). Energy and policy considerations for deep learning in nlp. *arXiv preprint arXiv:1906.02243*.

Taddy, M. (2019). The technological elements of artificial intelligence. *The Economics of Artificial Intelligence: An Agenda*, page 61.

Thoma, G. (2009). Striving for a large market: evidence from a general purpose technology in action. *Industrial and Corporate Change*, 18(1):107–138.

Trajtenberg, M. (2018). Ai as the next gpt: a political-economy perspective. Technical report, National Bureau of Economic Research.

Tubaro, P., Casilli, A. A., and Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1):2053951720919776.

van der Vleuten, E. (2009). Large technical systems. *A Companion to the Philosophy of Technology*, pages 218–222.

Van Roy, V., Vertesy, D., and Damioli, G. (2020). Ai and robotics innovation. *Handbook of Labor, Human Resources and Population Economics*, pages 1–35.

WIPO, W. (2019). Technology trends 2019: Artificial intelligence. *Geneva: World Intellectual Property Organization*.

Zhang, D., Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., Sellitto, M., et al. (2021). The ai index 2021 annual report. *arXiv preprint arXiv:2103.06312*.

Ángel Vázquez, M., Henarejos, P., Pérez-Neira, A. I., Grechi, E., Voight, A., Gil, J. C., Pappalardo, I., Credico, F. D., and Lancellotti, R. M. (2020). On the use of ai for satellite communications.

## Recent papers in the SPRU Working Paper Series:

**January**

2021.01. Exploring the links between research demand and supply: The case of Chagas. Valeria Arza and Agustina Colonna.

**December**

2020.20. Riskwork in the Construction of Heathrow Terminal 2. Rebecca Vine.

2020.19. The Origin of the Sharing Economy Meets the Legacy of Fractional Ownership. Francesco Pasimeni.

**November**

2020.18. Sustainability and Industrial Change: The Hindering Role of Complexity. Tommaso Ciarli and Karolina Safarzynska.

**October**

2020.17. Interplay of Policy Experimentation and Institutional Change in Transformative Policy Mixes: The Case of Mobility as a Service in Finland. Paula Kivimaa and Karoline S. Rogge.

**Suggested citation:**
Simone Vannuccini and Ekaterina Prytkova (2021). Artificial Intelligence's New Clothes? From General Purpose Technology to Large Technical System (SWPS), 2021-02. ISSN 2057-6668. Available at: www.sussex.ac.uk/spru/research/swps.

Science Policy Research Unit
University of Sussex, Falmer
Brighton BN1 9SL
United Kingdom

SPRU website: **www.sussex.ac.uk/business-school/spru**
SWPS website: **www.sussex.ac.uk/business-school/spru/research/working-papers**
Twitter: **@spru**

UNIVERSITY OF SUSSEX | BUSINESS SCHOOL