

SCIENCE POLICY RESEARCH UNIT

SPRU Working Paper Series

SWPS 2021-02 (May)

Artificial Intelligence' New Challenge? From General Purpose Technology to Large Technical Systems

Simone Vannuccini and Ekaterina Prytkova

Energy Policy

Sustainable
Development

Science and
Technology
Policy

Innovation
and Project
Management

Economics of
Innovation

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.

Editor

Roberto Camerani

Associate Editors

Karoline Rogge
Tim Foxon

Area

Energy Policy

Ben Martin
Ohid Yaqub

Science and Technology Policy

Andrew Stirling
Rob Byrne

Sustainable Development

Carlos Sato
Josh Siepel

Innovation and Project Management

Maria Savona
Alberto Marzucchi

Economic Policy

Editorial Assistant

Melina Galdos Frisancho

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/swps/

SSRN: www.ssrn.com/link/SPRU-RES.html

IDEAS: ideas.repec.org/workingpapers/sussex/

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 though 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

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 Gordon, 1997](#));

the types of impact we can expect from it, and guidance for the design of policies to govern it.

Approach . 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, GOF AI) and

²A 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 the four domains, each domain is characterized by a set of regularities and constraints. For instance, the domain of natural language processing is characterized by a set of regularities and constraints, such as the syntax, semantics, and pragmatics of a language. The domain of computer vision is characterized by a set of regularities and constraints, such as the geometry, topology, and color of objects in the environment. The domain of robotics is characterized by a set of regularities and constraints, such as the dynamics, kinematics, and sensors of a robot. The domain of artificial intelligence is characterized by a set of regularities and constraints, such as the logic, heuristics, and learning algorithms of an AI system.

(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 ([Angel Vazquez 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

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).

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 to capital. B31 h3(of)s(in2416-16.i3e8shisting)-by capital.sisation(tensiv)2ctions.ceoa(tensiv)2ctionsUSc mic aic

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](#)

	(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 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 standalone 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 derailed 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 brute-force 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 increasing 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

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

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 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 as stagnation, reconfiguration and decline ([Sovacool et al., 2018](#)). Furthermore, [Gokalp \(1992\)](#) 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 inefficiency" (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 Gokalp (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{

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 tailors 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 stepwise 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 or services, 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 content disclosure practices (the-ndhn [wc allesjustelo1 -16.(dec435 Td (((trading)-348g 0 0 1 F

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 domainsbeing 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 unmet input . This holds back or

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 problems¹⁶, 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 reverse 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




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

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) nite 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 di usion 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 unful lled: sustained commercialisation and growing competition among system builders make them race against each other, undertaking myopic steps in AI development leading to short{term payo . 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 identi ed 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 con ne the system, the mechanisms of control currently in practice, the distinctive style emerging, and nally 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 o ered 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 elds generates continuous feedback into the phases of innovation and development. Technology transfer has also accelerated with the increasing e orts of national, supranational and sub{national institutions to govern AI developments as the technology has

¹⁷The notion shares similarity with that of architectural innovation ([Henderson and Clark, 50](#)

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

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

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 underproduction 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 kickstarting 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

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.

In, Jones and Tonetti

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

nsiective to26292(dema's(e)-361((AI{usege)2629needsins.)-6(whatile)-361(demand2629is(e)-36foleplo),
pernority40009(of)-360(tseof)-363(solutions)-361(in)-360(tin)-36(along)-360{te.on

ts to high costs salong net258(w)28(ork)-768(e ects)-are

tslttshe developmentof2-297(the)-708(infrastructuts)-towards he 8(e cipmen)28(t)-crialisaties.thec(has
of proceocersandsmakfor8e ofir2(0)28(v)28opment p1(fdo-)]TJ 0 -16.435 Td [(noie)8292(ref8rof)y
rementadev

of algorithm 51.626 [(ln)-29assoo{econom1(ic)-29a(nst,ds)-409(ln)-298(8(e cipmen)28(t)-298((bstan)28(tl) cresting adfwution(migh)28(t)536(a(ash1(,)-46violresting)-29r(migh)28h1(,)-46dalagsting)5053(comp)72(meuati1(,

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

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 standalone 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

allocating 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.

Brodie, M. L. (1989). Future intelligent information systems: Ai and database technologies

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.

- 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. *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*.
- Kreuchau, 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.

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 paradigm *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 di usion 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 1(d)-3mis417(-ic)erolNe389(Expl2o4e389(Expl2g*

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

- 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.
- Angel Vázquez, M., Henarejos, P., Perez-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.

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.

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