



Tracing the evolution of service robotics: Insights from a topic modeling approach^{☆☆}

Ivan Savin^{*,a,b,c}, Ingrid Ott^{d,e}, Chris Konop^d

^a Institute of Environmental Science and Technology, Universitat Autònoma de Barcelona, Spain

^b Graduate School of Economics and Management, Ural Federal University, Yekaterinburg, Russian Federation

^c ICTA-ICP Building (Z), UAB Campus, Cerdanyola del Vallès, Barcelona, Spain

^d Chair for Economic Policy, Karlsruhe Institute of Technology, Germany

^e Kiel Institute for the World Economy, Kiel, Germany

ARTICLE INFO

JEL classification:

C11
C15
O33
O34

Keywords:

Knowledge diffusion
Latent dirichlet allocation
Networks
Patents
Topic matching

ABSTRACT

Taking robotic patents between 1977 and 2017 and building upon the topic modeling technique, we extract their latent topics, analyze how important these topics are over time, and how they are related to each other looking at how often they are recombined in the same patents. This allows us to differentiate between more and less important technological trends in robotics based on their stage of diffusion and position in the space of knowledge represented by a topic graph, where some topics appear isolated while others are highly interconnected. Furthermore, utilizing external reference texts that characterize service robots from a technical perspective, we propose and apply a novel approach to match the constructed topics to service robotics. The matching procedure is based on frequency and exclusivity of words overlapping between the patents and the reference texts. We identify around 20 topics belonging to service robotics. Our results corroborate earlier findings, but also provide novel insights on the content and stage of development of application areas in service robotics. With this study we contribute to a better understanding of the highly dynamic field of robotics as well as to new practices of utilizing the topic modeling approach, matching the resulting topics to external classifications and applying to them metrics from graph theory.

1. Introduction

Robots are increasingly supporting humans both at work and in their private life. While the use of industrial robots (IR) has a long standing tradition in the manufacturing industries, service robots (SR) are a more recent phenomenon. Latest advances in artificial intelligence and machine learning enable robots to sense and respond to their environments so that they can also be used outside secured production environments.

While IR still diffuse via intensified application in the manufacturing sector ('automation deepening'),¹ SR continuously capture new domains ('automation broadening'). Not always, but often, SR are mobile. Some of them are fully automatic or even autonomous.² Due to the importance of services in value creation and given the prevailing low level of automation in this field, future diffusion of SR is expected to have far-reaching implications for overall economic productivity. Due to the recent COVID-19 outbreak, this diffusion is experiencing a further boost

^{☆☆} Financial support from the Helmholtz Association (HIRG-0069) is gratefully acknowledged. Ivan Savin acknowledges support from the Russian Science Foundation [RSF grant number 19-18-00262]. This work has benefited from presentations at workshops in Kiel, Strasbourg, Karlsruhe, the EMAEE conference in Brighton and the EAEPE conference in Bilbao. All remaining shortcomings are our responsibility.

* Corresponding author.

E-mail address: Ivan.Savin@uab.cat (I. Savin).

¹ We use the term 'automation deepening' in the sense of intensified robot density, i.e. an increasing ratio of robots over employees. This is slightly different to Acemoglu and Restrepo (2018) who consider automation deepening as improvements of existing machinery.

² Following the international ISO standard 8373:2012, a robot is an actuated mechanism programmable in two or more axes, with a degree of autonomy, moving within its environment, to perform intended tasks. Given the current state and sensing, autonomy means the absence of human intervention by performing the task. The classification into industrial robot or service robot is done according to its intended application. The International Federation of Robotics (IFR, IFR (2018a)) bridges this standard to robot investment data broken down by application areas and industries. See Appendix A providing technical definitions of robots, robot systems, and the most important industries in which IR are used.

to facilitate more physical distancing in healthcare, logistics, tourism and other spheres (Chen et al., 2020; Yang et al., 2020; Zeng et al., 2020).

In this paper we aim to obtain a comprehensive understanding service robotics and of the associated technologies enabling their evolution. To this end, we use patent data, the most complete description of technological development. In line with the recombinant growth of knowledge, patent texts typically are composed of many technologies (Youn et al., 2015). To decompose these complex documents into distinct technologies, we take advantage of the topic modeling technique developed at the intersection of machine learning (ML) and natural language processing (NLP). In particular, we use not just their short abstracts that are too superficial to get an understanding of technologies incorporated in the patent, but full patent descriptions – or more precisely their non-technical summaries – in the robotic industry between 1977 and 2017. Then we extract the latent topics capturing different technologies used in patents, analyze how important those topics are over time, and how they are related to each other through their co-occurrence in patents.

In doing so we contribute to the existing literature in several ways. First, we apply modern methods of NLP and ML to exploit fine grained technological information included in unstructured textual data of patent documents and identify the optimal number of topics using established criteria of perplexity, exclusivity and coherence. We end up with 380 topics, carry out robustness checks for 190 topics and analyze how importance of those topics was changing over time. Second, we develop and apply a novel method for matching topics to SR based on an external text corpus provided by IFR that classifies SR technologies into 16 application areas and 49 sub-areas. We use results of this exercise to classify topics to different SR areas. Third, based on the results of our textual analysis we construct a complex graph using cosine similarity between the topics and identifying significant edges in this network. This approach allows us to overcome the popular practice of analyzing topics in isolation. Instead, we can trace robotic transformation from a system perspective: apply metrics from graph theory, understand the mutual relationship of discovered topics, distinguish between central (enabling) and peripheral (application) topics, discover communities of hardware- and software-oriented topics, and how those were changing over time. It is important to stress that all these steps are independent from structured patent data (e.g., patent classes, concordances) and expert bias. The entire process from topic identification until matching to SR is data-driven and thus is unbiased by subjective expert judgment. Another important strength of our approach is that apart from replicating results that could have been achieved by applying existing metrics (e.g. rising importance of medical robots, shift in popularity from hardware to software technologies), we are able to look inside the content of each particular topic (e.g. study technologies in surgery robotics) and their position in the complex space of knowledge comprised by robotic patents (e.g. how central they are and what topics they are connected to).

The remainder of this paper is organized as follows. Section 2 provides some background information on service robotics and topic modeling as a methodology to deal with patent data. Section 3 describes our data and methods. Section 4 presents the results. Section 5 discusses policy implications and Section 6 contains concluding remarks.

2. Background

2.1. The rise of service robotics

Research on automation took off with the introduction of mass-production manufacturing methods. Since then, automated machines have been continuously transformed to today's multipurpose industrial robots, i.e. robots that are able to adapt to a different application without alteration of the mechanical system. Until today, key drivers of robot diffusion have been *automation deepening*, i.e. an intensified use of

robots in already automated industries (as e.g. the automotive and the electrical industry) or the increased penetration of robots in countries with hitherto still only few industries being automated. Driven by increased calculation power, a decline in hardware costs, the use of lighter materials and technological progress in complementary technologies, such as cloud computing or artificial intelligence (AI), robot use continuously expands into new fields that until today have been characterized by almost complete absence of automation.

However, most of research on robots in economics has still the focus on IR, mostly associated with labor market implications, task perspective and addressing questions like the relationship between humans and machine being complementary or substitutive (e.g. Acemoglu and Restrepo (2019) or Graetz and Michaels (2018)), covering national perspectives (e.g. Dauth et al. (2018), Dauth et al. (2019) for Germany, Acemoglu et al. (2020) for France, Bessen et al. (2020) for the Netherlands). Agrawal et al. (2019) focus on prediction of labor market implications. A more historical dimension has been taken by Atack et al. (2019) or already Mansfield (1989), who analyze diffusion of robots in the US and Japan, and by Cheng et al. (2019) for China. Baldwin and Forslid (2020) are among the few economists that also explicitly relate their work to services and robots in a globalized context. Autor and Dorn (2013) point out that it is crucial to distinguish between service robots is still rare, especially in the scientific literature. Recent robot patent analysis can be found by Analytics (2018) or EPO (2017), while Goeldner et al. (2015) are among the first who analyze patenting activity in a specific application field of SR, namely care.

Modern robots are flexible, easy to operate and becoming able to navigate autonomously, even in unstructured environments. As a consequence, aside from using robots within clearly defined environments such as factories, provision of services is increasingly becoming automated. The associated spread of robot application may be labeled as *automation broadening* providing potential for huge productivity gains from automation also beyond the manufacturing industries.

According to the ISO classification, SR perform useful tasks for humans or equipment excluding industrial applications.³ SR are further differentiated into two types, namely those for *private use* that are operated by a lay person and those for *professional use* that are usually operated by a professional or properly trained operator. Due to the multitude of forms and structures as well as application areas of SR, it is not easy to delimit them from IR. Robots in logistics are a prominent example of such an unclear assignment. They are used in non-manufacturing environments, such as logistic centers, hospitals or warehouses but also to transport parts within factories.⁴

The sketched technological transformation is being mirrored in the work of the IFR's statistical department. It bridges robot classification from the ISO standard and robot markets by collecting data on worldwide robot investments, sales and stocks differentiated by IR and SR; the latter ones split up into robots for domestic and for professional use. For the heterogeneous domain of SR, the IFR provides a detailed technology breakdown covering 16 areas and 49 sub-areas (Table 1).⁵ In 2017 the most important markets of SR have been in the fields logistic systems, defense applications, public relation robots, field robots (especially milking), powered human exoskeletons and medical robots (IFR, 2018b). According to IFR (2019), sales volumes of logistic robots have been the key driver of the SR markets also in 2018 followed by SR

³ See ISO 8371:2012, 2.11 (private use; synonyms are personal or domestic use) and 2.13. (professional use).

⁴ In 2018, 7'700 units of logistic robots have been used in manufacturing while 103'000 units have been utilized outside of factories, e.g. in warehouses, logistic centers and hospitals (IFR, 2019).

⁵ The IFR provides information differentiating between IR and SR and categorizing by application area, industrial branches, robot types or geographical region, and across time.




Fig. 1. Investment dynamics for IR and for selected SR application areas. Data is normalized so that the investment level in 2016 equals one. In both panels the red line refers to the level of one. Data for IR dynamics (until 2018) based on IFR data base; data on SR dynamics based on IFR (2018b) and IFR (2019). Starting from 2018 the investment data is predicted by IFR. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

applied for inspection and demolition tasks or medical service, most of them being surgery robots. Considerable markets exist in field robotics (e.g. milking), and big potential is seen in the application of robots for plowing (agriculture). Another fast growing market is public relations robots which are used to provide information in public spaces or shops thereby also increasingly utilizing humanoid robots. SR for private use are dominated by sales of lawn mowers or floor- and window-cleaning robots, together with robotic toys and games. Technological advancements in robot mobility also drive the adoption of robots in the (still small) market for elderly and handicap assistance. Finally, SR for domestic use are reported separately. Their *unit value* is only a fraction of that of the many types used for professional use. Since they are produced for mass markets, they follow different pricing and marketing channels.

The above summary illustrates the strong heterogeneity not only between IR and SR, but especially also within SR applications. They also differ with regard to unit price,⁶ life span⁷ and investment dynamics.

While IFR reports reliable data on IR since 1993 onward, comparable time series are not available for SR. A central reason for this is incomplete data due to the high fluctuation of providers in this dynamic market segment. Another problem is that the reported data is difficult to compare over time. Since the focus of our analysis is on the transformation of robot technology, data on absolute investments play only a minor role. Instead, the dynamics of the specific SR areas, their evolution over time and their embedding within the robot technology space are important. In the following, we present the data on SR from IFR (IFR, 2018b; 2019). Recent investment dynamics as well as forecasts for 2019–2022 is approximated based on the IFR's annual information on sales. visualizes this by setting the unit sales of the respective robot area in the year 2016 to one and showing the dynamics of factual and predicted data. Fig. 1(a) reports IR investment dynamics beginning in 1993. Since it has a longer history and higher levels of installed units, their development is less dynamic and even stagnates in 2018. Investment dynamics of selected SR areas are displayed in Fig. 1(b) demonstrating SR growing many times faster than IR with the leading areas being logistics and medicine.

⁶ Medical and underwater robots sometimes cost several hundred thousand USD (IFR, 2018b), while toy robots often only a few hundred USD

⁷ For industrial robots, the average duration of use is about 11 years, underwater SR are utilized up to 10 years, while defense robots may have a life cycle of one single operation.

To better understand the future development of robotic technologies, we take a precise look at patenting as an early indicator of future investment in the respective area. In robotics, there is a strong correlation between patenting and investments (Fig. 2). This correlation pattern also holds if one zooms into major geographical regions or robot applications (like handling or soldering). It is also worthwhile to stress that robot patenting has been less susceptible than investment during and in the aftermath of the global financial crisis. These observations together with the aforementioned description of the transformation of robotic technologies is the starting point of the subsequent empirical analysis of patents with a special - though not exclusive - focus on emergent SR.

2.2. Patent data analysis and use of topic modeling

Patent data is widely considered as the most complete description of innovations. Patents cover a multitude of technical fields over a long period of time, and thus reflect structural changes in those technologies. There are specific criteria for an invention to be patented, which creates an objective standard as to what counts as an invention. Information on patent citations has been long used in economics as an indicator for patent quality (Trajtenberg, 1990), knowledge diffusion (Jaffe et al., 1993) and its obsolescence (Jaffe and Trajtenberg, 1996).

The use of patent citation data, however, has been always problematic due to different practices of patent citation across patent offices with citation records being often incomplete (Michel and Bettels, 2001) and, most important for the reasoning of the paper at hand, the fact that patent examiners may have added extra citations not actually relevant for the inventions (see Alcácer and Gittelman (2006) for USPTO evidence). This bears the risk of distortions (expert bias) if one wants to analyze the development of technology over time. Furthermore, citations are usually made on the basis of legal considerations rather than from a technological perspective. Besides, patent classifications change over time (Lafond and Kim, 2019) and are also hard to compare between different patent offices. All this is challenging the use of structured information from the patent data, which is considered as the most comprehensive and accurate description on knowledge flows. What is most striking is that a large body of information unstructured textual description contained in patents) has been long disregarded in economic research.

With the rising power of modern computers and availability of the great amount of data, however, our choice of instruments to extract




Fig. 2. Robotic patents granted at USPTO (1077–2016) and robot investment worldwide (1993–2016).




Fig. 3. Length of patent non-technical summaries.

information from the textual data is rapidly growing. In this study we apply the so-called topic modeling (TM) approach to gain new insights about knowledge structure and dynamics in the robotic industry. TM is a clustering approach for textual data aimed to identify meaningful topics in text data, analyze trends in topics, (re)classify and annotate documents (Blei, 2012; De Battisti et al., 2015). TM posits the idea that patents represent a combination of topics, where each topic is a probability distribution over a fixed vocabulary. While patent texts are observed, the topics and their distributions are not and treated as latent variables. TM, thus, uncovers topics latent in a collection of patents and identifies which composition of them best accounts for each patent. The advantage of TM over, e.g., keyword analysis is that i) it allows words to have different meanings depending on their contexts; and ii) it is data-driven: one does not need to specify topics *a priori* but generates them from the data. More details on topic modeling related to our data set are presented and discussed in Section 3.3.

TM has been recently applied to patent data in a number of studies. It has been used for patent (re)classification into product and technology sub-classes to later explore technological convergence and geography of innovation in the photovoltaic technology between the US states (Venugopalan and Rai, 2015); for identification of emerging topics among triadic patent families (patented in the US, EU and Japan, Lee et al. (2015)); for detection of pioneering patents introducing new topics (Kaplan and Vakili, 2015); and for prediction trends in patent topics (Chen et al., 2017; Suominen et al., 2017). To the best of our knowledge,

robotic industry has never been well studied with respect to the topics prevailing in the related patents.⁸

Apart from patents, TM has been widely applied to other type of textual information. Many studies focus on scientific literature published either in a specific peer-reviewed journal across many themes (De Battisti et al., 2015; Griffith and Steyvers, 2004; Lüdering and Winker, 2016; Savin and van den Bergh, 2021) or all economic articles stored in a given database (such as JSTOR, Ambrosino et al. (2018)). Furthermore, some studies focus specifically on literature published on the theme of information security (Chang, 2016) or bioinformatics (Liu et al., 2016). While some of these studies look on abstract only (like De Battisti et al. (2015)), others take full texts of the academic papers into analysis (Ambrosino et al., 2018). Another popular field of application for TM is news articles. Those can be either specialized financial news (taken, e.g., from Dow Jones Newswires Archive (Larsen and Thorsrud, 2019) or financial analyst reports (Huang et al., 2017), policy statements and website articles related to climate change (Farrell, 2016) or publications from social media like Twitter (see, e.g., Chae and Park (2018)). Finally, topic modeling has been recently applied to survey open-ended questions (Roberts et al., 2014; Savin et al., 2021; 2020; Tvinnereim et al., 2017). This illustrates the generality of the approach that can be applied to very different type of data in terms of size and content.

3. Data and methodology

3.1. The robotics patent data set

We focus on robotics utility patents granted at the United States Patent and Trademark Office (USPTO), where we take the first filed patent of a family as representative.⁹ Our search strategy is based on having the truncated keyword **robot** in the title or abstract of the patent¹⁰ or being classified in one of the CPC classes concordant to the

⁸ The only exception perhaps is by Kim et al. (2016) focusing on a small fraction of robotic patents devoted to humanoid robots using less than 1000 patents from USPTO and analyzing their titles and abstracts only.

⁹ A family of patents refers to a group of patents that are issued in different countries for the same invention to obtain patent protection. We refrain from multiple counts of the same invention within the patent families since the size of the family does not affect the technical facets of the invention. If we were to include every family patent in the analysis, this would distort the text corpus in favor of the larger families.

¹⁰ We initially applied the same search strategy on the full texts of patents at the USPTO database. However, when checking the result it became obvious that this strategy yielded too many false positives.

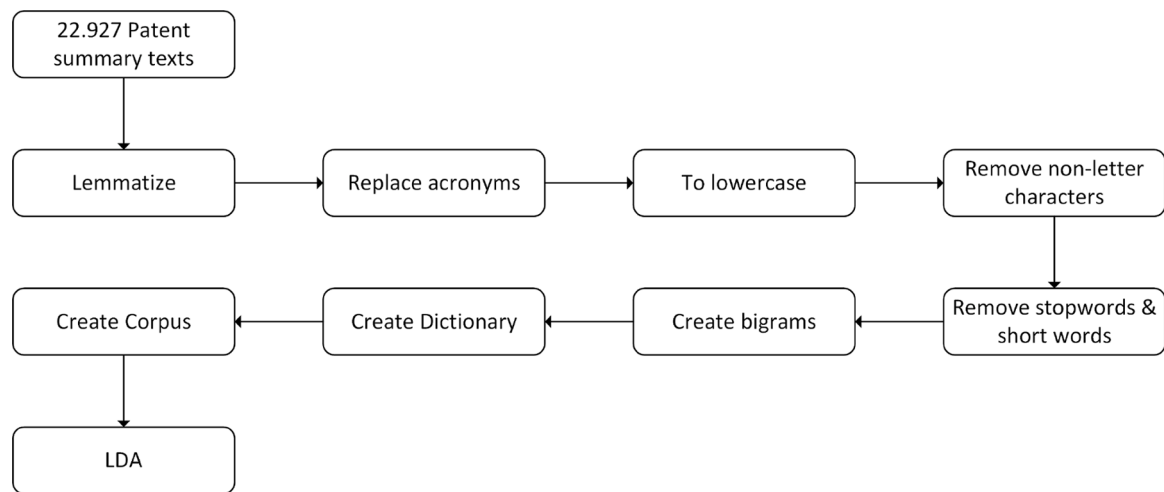


Fig. 4. Data cleaning procedure .

USPC class 901 (robots).¹¹ This identification strategy resulted in 22'927 patents for the period 1977–2017 (see Fig. 2; black line for the evolution of patents over time).

Patent texts consist of several parts which can be considered for textual analysis: title, abstract (typically less than 150 words), patent description and patent claim. While title and abstract are too short to get any comprehensive understanding of technologies incorporated in the patent, the claims concentrate on the differences of the technical novelty compared to the prior art. Claim texts are thus also with limits suitable to reflect the content of new technological knowledge contained in the patents. The patent description includes a non-technical summary as well as a technical description. For our analysis, we skip the latter and concentrate on the non-technical summary. This has the advantage that we avoid parts of text with too specific language and formula description that is less suitable for textual classification. On average, the length of the non-technical patent summary descriptions is 667 words (see Fig. 3 for descriptive statistics on the data). In total, 15.3 million words are contained in non-technical summaries of our patent sample.¹²

Besides, each patent includes some structured data including classification to one or several patent classes. Existing concordances make it possible based on the evaluation of the CPC classes of patent specifications to assign the patents to industrial classification, e.g. by means of NACE or ISIC concordance classifications Eurostat (2014). Similarly, patents can be assigned to technology fields (by means of WIPO technology area concordance matrices WIPO (2008) with a strong bias towards manufacturing industries. Based on this information it is possible, for example, to visualize the development of medical instruments or the increasing importance of ICT components by exploiting structured information of our robotic patent dataset (see Fig. A12 in A.2). Applying both concordances illustrates a shift over time from mechanical components to computers, an increase in medical instruments, a decline in machine tools while measuring, testing and navigation remains rather stable. We will later illustrate that our approach is capable to replicate this pattern.

¹¹ See <https://www.uspto.gov/web/patents/classification/cpc/pdf/us901tocpc.pdf>. A full list of concordant CPC classes can be found at: <https://www.uspto.gov/web/patents/classification/cpc/html/us901tocpc.html#statTable>.

¹² We also conducted our analysis on the full descriptions (both non-technical and technical description) as well as the technical descriptions only. With the same number of topics chosen the results are fairly similar concerning the topic content and are available upon request. However, as technical summaries contain more formulas and other scientific notation, by concentrating solely on non-technical summaries we avoid supplying our NLP analysis with non-textual data that is hard to clean automatically.

3.2. Patent data pre-processing

The first step of our pre-processing pipeline is lemmatization, where inflected word forms are transformed to their dictionary form. Contrary to stemming the words (i.e. cutting words to their word stem, or root form), lemmatization is done by identifying the intended part of speech and meaning of a word in a sentence. We tried both approaches and while both of them have limitations, topic keywords generated after stemming are harder to interpret as they usually contain only parts of the words. Since multiple words with different meaning may be grouped into the same stem, using lemmatization allows us to preserve interpretability of our results. In particular, we use WordNetLemmatizer from the Natural Language Toolkit (NLTK) in the Python programming language.

Second, after lemmatization, we replace acronyms by their original notation. This step is very important, as patents, similar to academic literature, contain a lot of acronyms introduced once at the beginning of patent description and used consistently in the document. Overall we find 1502 distinct acronyms, resulting in a total of 18'628 replacements. By identifying and replacing those acronyms with their actual meaning we solve several problems:¹³ (i) we reduce the chance that the same acronym used in different documents and actually standing for different word combinations will be recognized by our approach as the same word. Thus, although we find 1502 acronyms, there are only 1028 unique acronyms, meaning 474 acronyms have the same sequence of characters but a different meaning behind them. This implies that without replacing acronyms we would have biased our results by not distinguishing the terms with non-unique acronyms. (ii) we considerably increase our ability to understand the formed topics later on as instead of an acronym we can see the full expression.

Third, this is followed by transforming all characters to lowercase and then removing stopwords (i.e. and, or, the) as well as all non-characters.¹⁴ This step is standard for NLP (see, e.g., Grün and Hornik (2011)). Note that we do that after replacing acronyms and lemmatizing words to minimize the amount of information that may be lost (e.g.,

¹³ The replacement is done by an n-gram with the words being connected in one with "." symbol.

¹⁴ In particular, the NLTK's list of English stopwords has been used extended with a set of custom stopwords we created while testing different LDA models (see Section B.1 for more details).

short acronym of two letters only).¹⁵

Fourth, we create bi-grams out of words commonly occurring together. Bi-grams are created using gensim phrase module in Python. For each bi-gram, the normalized pointwise mutual information (NMPI) score has been applied (see Bouma (2009)). In simple words, the NMPI score measures how often any two words appear together versus how often they appear in text separately and forms out of those appearing predominantly together a bi-gram with a "_" symbol (see Appendix B for the exact formula of the score and its explanation.). Specifically, a 0.5 threshold value to form bi-grams is used. The flow of the pre-processing steps on our textual data is summarized in Fig. 4.

3.3. Topic modeling

While there have been many algorithms developed for topic modeling (for a recent overview see Liu et al. (2016)), in the following we concentrate on the most commonly used approach known as latent Dirichlet allocation (LDA, Blei et al. (2003)). LDA as a probabilistic modeling approach assumes that each word in a document is generated in two step process. First, assuming that each document has its own topic distribution, a topic is randomly drawn from that distribution (the Dirichlet distribution). Second, assuming that each topic has its own word distribution, a word is randomly drawn from this word distribution of the topic selected in the first step. A document is essentially a result of multiple repetitions of the statistical model consisting of those two steps the generative process defined by a joint probability distribution over the observed (documents and words) and hidden (topics) variables. This statistical model reflects the intuition that documents exhibit multiple topics, where each document exhibits the topics in different proportion. The LDA algorithm discovers the topic distribution for each document and the word distribution of each topic iteratively, by fitting this two-step procedure to the observed documents until it finds the best set of variables that describe the topic and word distributions. Effectively, LDA as a Bayesian model computes the conditional distribution of the hidden topics given observed data identifying the posterior distribution of the latent topics in a collection of documents.

Similar to cluster and principal component analysis, LDA reduces the dimensionality of linguistic data from words to topics based on the co-location of words in a collection of documents to infer the underlying topics in those texts and the weight of each topic in each individual document. According to Blei (2012, p. 79): "This can be thought of as 'reversing' the generative process - what is the hidden structure that likely generated the observed collection". For a more formal description of LDA, see Appendix B.

The number of topics in the model affects the interpretability of its results. Setting the number too low can result in topics that are too broad and ambiguous. Conversely, setting the number too high may introduce uninterpretable topics that pick out idiosyncratic word combinations (Griffith and Steyvers, 2004). While one can follow computational linguistic literature and calculate perplexity score (fitness) of the LDA model based on different numbers of topics (Blei et al., 2003), most of the time this number is chosen ad hoc and relatively small to maintain interpretability of results. It is by all means easier both to label and overview 10 and not 100 topics.

In our study we will choose the number of topics that aims to optimize not just perplexity score of the model, but also semantic coherence and exclusivity of the topics. In simple words, the perplexity score is goodness of the LDA model trained on one part of the data to predict the word distribution on the other part of it. Semantic coherence quantifies

the extent to which frequent words from the same topic tend to appear in the same patents. Exclusivity analyzes whether popular words from one topic have low likelihood in all other topics. Thus, we follow Roberts et al. (2014) in arguing that semantically interpretable topics should consist of words that tend to co-occur within documents, and that their top keywords are unlikely to overlap with keywords from other topics. More details on the different metrics used to determine the optimal number of topics are provided in Appendix B.

Applying the topic modeling approach on the patent data results in two matrices. The first is the matrix of probabilities observing a word given the topic, while the second is the matrix of topic prevalences in each of the documents. The information in these two matrices is the basis for labeling topics, visualizing them and conducting further analysis such as matching the formed topics to existing SR classification and analyzing mutual interdependence between topics by means of network analysis. These two steps of analysis we address in the following subsections while subsections 4.2 and 4.3 provide the corresponding results.

To illustrate the content of each topic and evolution of their importance over time, we will henceforth use *word clouds* and *diffusion curves*. Word clouds represent the 30 most frequent words given the topic with font size capturing how likely these words are given the topic. The lightness of the color in turn reflects how exclusive this word is compared to all other topics: the lighter the color, the more exclusive is the word. Diffusion curves represent the stage of technology adoption captured by our topics in the corpus of robotic patents over the period 1977–2017. To produce such diffusion curves, we follow Lenz and Winker (2020) in quantifying the probability that a given topic appears in the corpus of patent texts for each year ensuring that that these probabilities for any period sum up to one. In addition, we smooth the curves by estimating such probabilities not just for a single year but for a five-year time interval around a given period (see Appendix B for more details). Finally, to ease reading these plots, we classify the curves into 'rising', 'falling' or 'in-between' with green, red and blue color, respectively, depending whether the topic between 1977 and 2017 was predominantly rising or falling in its prevalence, or none of the two (see Appendix B for details).

3.4. Topic matching

While methods to form topic models have seen a rapid development over the last two decades, labeling and interpreting topics remains largely an ad hoc procedure. Most studies try to provide a concise label consisting of frequent and exclusive words summarizing the essence of the topic. Apart from looking on single terms associated with the topics, some papers in addition take illustrative documents where the prevalence¹⁶ of respective topics is highest to demonstrate how topic labels fit in the context (see Roberts et al. (2014); Savin et al. (2020); Tvinnereim and Fløttum (2015)).

In a recent survey, Boyd-Graber et al. (2017) distinguish between labeling methods that only use internal information from the topic model against those that also use external knowledge sources. Labeling with exclusively internal information looks for phrases with high topic prevalence that well summarize the documents making them good candidates for labels. Labeling with external knowledge sources either aims to weight words in a topic as prospective labels putting more weight on words that are hypernyms¹⁷ (which is assessed through external word library) and that co-occur often with other words from that topic, or try to retrieve labels of the documents underlying the topics (if available), and form topics in line with those labels. Applications of such automatic topic labeling are very scarce (see e.g. Newman

¹⁵ Note that we decided against applying term-frequency inverse document frequency filter, as we could delete too many relevant words necessary later for our topic matching. Since we used for our simulations the high-performance computing cluster BwUniCluster https://wiki.bwhpc.de/e/Category:BwUniCluster_2.0, this was feasible in terms of computational time.

¹⁶ Related to our patent data set, topic prevalence refers to the degree (between 0 and 1) to which a patent document belongs to the respective topic.

¹⁷ Hypernyms are words which meaning includes a group of other words, e.g. related to the word dog a hypernym would be pet.




Fig. 5. Perplexity, coherence and exclusivity for LDA models with different number of topics.

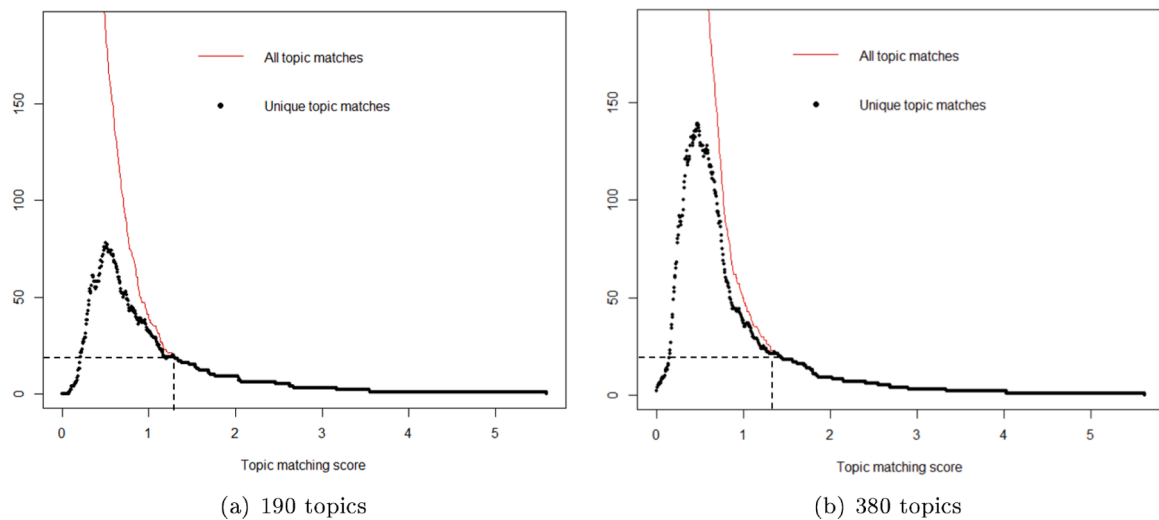


Fig. 6. Deriving a threshold value for topic matching. *Note:* The red line shows how many topics are matched in total (i.e. including the same topic being matched to two or more SR sub-areas) given the threshold value, while the black one shows how many topics are matched uniquely to one SR sub-area as listed in Table 1. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

et al. (2010)) demonstrating a large room for improvement in methods that can objectively label topics or match them to some external classification.

One of the major contributions of the paper at hand is the development of a matching algorithm that allows for automatic identification of SR topics, i.e. independent of subjective judgment (i.e. domain knowledge of experts), specifically at the level of SR sub-areas in Table 1. To the best of our knowledge, we are the first to propose such a method for topic assignment to a specific technology field based on patent data.

For this, identification of an appropriate reference text distinguishing between SR areas based on textual information from the topics resulting from LDA is essential. In our case we rely on textual description of SR technologies as provided by the IFR SR report 2018 (IFR, 2018b, chap 3, pp. 48–270) containing a detailed textual description of SR technologies at the sub-area level).¹⁸

More precisely, we have two different textual sources relevant for

matching, one being the lists of words associated to each topic with different probabilities¹⁹ and the other one being the descriptions of SR areas from IFR, which serve as reference text. Then we look for words that appear both, in the topics and in the SR description, labeling them as 'candidate words'. Based on these words we calculate for each topic a 'topic matching score' for each of the existing 49 SR sub-areas capturing the extent to which they overlap. These topic matching scores tend to be higher if words from a given topic have a higher frequency (i.e. probability belonging to this topic) and exclusivity (i.e. low probability of appearing in any other generated topics) and, symmetrically, appear more often in the description of the given SR sub-area and less often in other SR sub-areas. Furthermore, since we use only SR descriptions in the reference texts (and not for example, descriptions of other technological fields such as IR and beyond, we additionally reduce the topic matching scores for those words, which are more common in the

¹⁸ We also tried to apply the matching procedure on the level of 16 SR areas. We prefer, however, to look on the less aggregated classification level to attribute topics more precisely. Knowing the SR sub-area, it is straightforward to identify the corresponding area.

¹⁹ As was mentioned earlier, the main output of LDA is a matrix of size the number of topics times the number of words in the corpus, where each word is attributed to every topic with different probabilities.

Table 1
SR applications: 16 areas (bold) and their 49 sub-areas as defined in IFR (2018b).

Service Robotics	
for private use	for professional use
Robots for domestic task	Field robotics
Robot companions, assistants, humanoids	Agriculture
Vacuuming, floor cleaning	Milking robots and livestock robotics
Window cleaning	Mining systems
Lawn-mowing	Space robots
Pool cleaning	
Entertainment robots	Professional cleaning
Toys and hobby robots	Floor cleaning
	Window and wall cleaning (incl. wall-climbing robots)
Multimedia robots	Tank, tube and pipe cleaning
Education and research	Hull cleaning (aircraft, vehicles, ships etc.)
	Other cleaning tasks
Elderly and handicap assistance	Construction and demolition
Robotized wheelchairs	Nuclear demolition and dismantling
Personal aids and assistive devices	
	Building construction
Home security and surveillance	Heavy/civil construction
Home security and surveillance	Other construction systems (road construction)
	Logistic systems
	Automated Guided Vehicles (AGVs) in manufacturing environments
	AGVs in non-manufacturing environments (indoor)
	Cargo handling, outdoor logistics
	Personal transportation
	Inspection and maintenance systems
	Facilities and plants
	Tank, tubes, pipes and sewers
	Other inspection systems (inspection robots for nuclear plants)
	Medical robotics
	Diagnostic systems
	Robot-assisted surgery and therapy
	Rehabilitation systems
	Other medical robots
	Rescue and security applications
	Fire- and disaster-fighting robots
	Surveillance/security robots
	Other surveillance and security robots
	Defense applications
	Demining
	Unmanned aerial vehicles (defense applications)
	Unmanned ground-based vehicles
	Unmanned Underwater Systems
	Underwater systems (civil / general use)
	Underwater systems (civil / general use)
	Powered human exoskeletons
	Powered human exoskeletons
	Mobile platforms in general use
	Mobile platforms in general use
	Public relations and joy rides
	Hotel and restaurant
	Guidance
	Marketing
	Robot joy rides

scientific literature of contemporary American English.²⁰ As a result, each topic belongs to each SR sub-area with a different matching score.²¹ In the final step, we derive a threshold value for these matching

scores that allows us to uniquely classify the largest number of topics to SR sub-areas. This last step is data-driven, i.e. instead of choosing the threshold externally we let the data decide what should be the value of the matching score to assign a topic to any SR sub-area. In Appendix B we explain all steps of the matching algorithm in detail, while Section 4 presents the results of this approach.

3.5. Network space of topics

Having classified the patent descriptions into topics, we can assess to what extent those topics are related to each other, i.e. how often these topics appear in the same patent descriptions. The underlying assumption is that the more the knowledge on technologies captured by the two topics are used jointly in the same patents, the more they are interdependent. An example can be navigation of autonomous guided vehicles based on the method of scanning the external environment. Capturing these relations between topics is important as it helps to look on the topics not merely in isolation which is not adequate to understand distinct technologies and to arrive at an overall picture of the knowledge space underlying the robotics industry. Therefore, the approach presented below is the second methodological contribution of the present paper to the literature applying topic modeling to patent data.

For this purpose, we define two topics as connected based on their *cosine similarity*, i.e. co-occurrence of those topics in the same patents (see Appendix B for a formal description of the measure). Measuring this similarity between any pair of topics we obtain a symmetric matrix with ones on the main diagonal as cosine similarity of two identical vectors is one, and all other values bounded in [0,1]. These values capture the strength of relation between any pair of topics and can be interpreted as corresponding edge weights in an undirected graph between these topics. In particular, a high weight implies that the two topics appear in many patents with a high prevalence, and they are strongly linked. Since one patent can contain small parts of text belonging to many topics, we get a virtually fully connected graph, where many edges have relatively low weight. To simplify its analysis and visualization, we assess significance of each particular edge in this weighted graph. To do so, we follow Saracco et al. (2015) in constructing an appropriate *null model* – its randomly generated counterpart – which displays on average the same degree distribution (diversification of patents) and the same ubiquity (weight distribution of each topic) 1000 times.²² Comparing the empirically observed weighted graph with their randomly generated counterparts we preserve only those links which weight surpasses the 95% threshold, i.e. they fall in the 5% of most outstanding edge weights which could have been observed given the underlying data. The 5% threshold is taken as most conventional significance level observed in the empirical literature. Note at this point that by deleting 95% of links and transforming our network from weighted to unweighted one, we reduce the density of our topic networks to 0.05, i.e. only 5% of all possible links are present in our network. This, however, does not preclude us from analyzing other important characteristics of the topic network such as *number of components* capturing how many topics are isolated from other topics or *network centrality* accounting for the concentration of edges on few topics (degree centralization) and the dependence on topics that connect many other topics (betweenness centralization) proposed by Freeman (1979). It is important to remember, however, that the absence of an edge between any pair of topics in our henceforth analysis does not imply that they never appear together in any patent, but that the extent to which they co-occur does not meet the bar to be considered as significant.

²⁰ To this end, the database composed by Mark Davies (<https://www.wordfrequency.info/>) has been used.

²¹ Logically, if there was no candidate word appearing in the topic with a positive probability and in the specific SR sub-area, the corresponding matching score equals zero.

²² The same approach for testing link significance has been applied, among others, by Napolitano et al. (2018) and Pugliese et al. (2019).

Table 2
Topics from LDA model with $K=380$ matched to SR sub-areas.

Topic	Topic terms (30 most relevant)	Matched SR sub-area
Field Robotics		
377	platen, flat, abrasive, crop, agricultural, harvest, harvester, tractor, abrade, lapping, tree, harvesting, label, lap, fertilizer, farm, abrading, granite, fruit, agriculture, barcode_label, edging, farmer, queue_manager, floating_platen, hieroglyph, flat_surfaced, flatness, edger, contractor,	Agriculture
86	teat_cup, milk, animal, milking, teat, cow, milk_animal, harness, milking_animal, teat_animal, wire_harness, milking_parlor, udder, animal_milk, luggage, fodder, milk_parlor, udder_animal, disinfect, slotted, hind_leg, animal_udder, jumper, foremilk, occupation, beveling, facet, formboard, compartment, milker,	Milking robots and livestock robotics
330	turret, conversion, last_teat, stepper_motor, light_curtain, deactivation_threshold, dairy_animal, strain_relief, umbilical_cord, milking, averaging_period, duration, garage_door, exploration, mending, position_sensor_print_circuit_board, staircase, intensifier, nonlinear, recirculating, nonlinear_exploration, pause, unprogrammed, bipedalism, historical, dependence, gunshot, sclera, touchup, obstruction,	Milking robots and livestock robotics
113	shift, traction, bucket, orbit, momentum, recite_claim, zmp, spacecraft, shifting, leg_hopping, upstanding, hst, midpoint, circumscribed, ferrule, stable, foamed_seal, insole, forwardly, shiftable, outwardly_shifts, mule, virus_strain, infectious_bacterial, wearable_assistive, backwardly_shift, inscribed_cylinder, outwardly_flexed, backwardly_flex, time_of_flight_diffraction,	Space robots
37	seed, sucker, hood, mat, baseplate, nutrient, tissue_explants, soil, seedling, grit, firefighting, curved_backplane, handheld, greenhouse, plant, loaded_filament, podium, seeding, tissue_explant, mature_plant, debridement, chimney, planting, leaf, sucker_revolver, harvest_crop, public_speaking, germination, presentment, suspendably_translatable,	Agriculture
260	course, unmanned_vehicle, truck, unmanned_dump, dump, dump_truck, topographical, foremilk, earth_sand, wherefore, partial, contraband, haul, travel, survey, anonymous, wherewith, twice_diffacted, loader, presence_contraband, conflict, wait, manned, loading, idea_underlie, quarry, dumping, site, topological_map, lashing,	Mining systems
Medical robotics		
83	fiducial, jog_feed, canister, stereotactic, cockpit, transformable_toy, toy_doll, craft, skull, shin, neurosurgery, preoperative, astronaut, fiducials, fixable, sanding, sand, neurosurgical, anatomical, tiptoe, clinch_nut, aquatic, predrilled_hole, platter, satellite_platter, mri, intraoperative, raster, microsurgery, intracranial,	Robot-assisted surgery and therapy
6 ↑	catheter, elongated, distal_end, medical, sled, elongate, handle, proximal_end, sheath, disposable, catheter_sheath, sterile_barrier, nose_cone, sterile, steerable_catheter, patient, ablation, knob, ablation_catheter, introducer, glass_crucible, sterility, catheter_introducer, electrophysiology_catheter, remote, introducer_sheath, mapping_ablation, electrophysiology, remotely, effective,	Robot-assisted surgery and therapy
191 ↑	surgical, surgeon, surgical_instrument, surgery, minimally_invasive, surgical_procedure, surgical_site, incision, patient, medical, instrument, medical_procedure, endoscope, entry, intuitive_surgical, robotic, console, system, tissue, procedure, vinci_surgical, trauma, endoscopic, hospital_stay, perform, inc_sunnyvale, robotically, stereoscopic, orthopedic, recovery,	Robot-assisted surgery and therapy
280 ↑	tissue, organ, accessory, forceps, endoscope, laparoscopic, capsule, endoscopic, trocar, laparoscope, blood_vessel, snake, laparoscopic_surgery, endoscopy, laparoscopy, minimally_invasive_surgical, abdomen, abdominal_cavity, continuum, vivo, abdominal, cavity, terminus, minimally_invasive, intracorporeal, cannula, connected, discomfort, clip_applier, invasive_surgical,	Robot-assisted surgery and therapy
167	radiation, tumor, dose, patient, diagnostic, couch, ray, therapeutic, lesion, imaging, radiation_therapy, cancer, radiotherapy, therapeutic_radiation, treatment, therapy, ambulatory, sighting, radiosurgery, physician, biopsy, compute_tomography, magnetic_resonance, shooter, rays, computed_tomography, collimator, medical, proton, linac,	Robot-assisted surgery and therapy
311 ↑	simulation, user, training, simulate, simulated, feedback, simulator, physical, haptic, virtual_reality, graphical, computer, mouse, interaction, joystick, trainee, intuitive, real, feel, sensation, train, interface, realistic, interact, amusement, haptic_feedback, experience, environment, graphic, assistive,	Rehabilitation systems
308 ↑	prosthetic, prosthesis, ankle, artificial, wearer, amputee, knee, hip, foot, gait, flexion, limb, residual_limb, powered, muscle, prosthetic_knee, toe, automatic_guided_vehicle, extremity, ankle_joint, amputation, flexion_extension, prosthetic_orthotic, damper, myoelectric, orthosis, heel, knee_joint, thigh, heel_strike,	Rehabilitation systems
Robots for domestic task		
130 ↑	rack, boundary, mower, lawnmower, lawn, grass, mow, mowing, saw_blade, helical_spring, mow_lawn, robotic_welding_assembly_device, root, geographic, traversal_launch, urls, mowable, traversal, substage_await, nursery_school, steeple, cubic, substage, substage_empty, photo_interrupter, rotor_steeple, young_child, unburned, angioplastic, playmate,	Lawn-mowing
284	clean, dry, wash, cleaning, rinse, bath, wet, washing, drying, chemical, foreign, residue, cleaned, solution, jet, liquid, dirty, pellicle_paste, impurity, immerse, remove, particle, towel, ipa, spray, diced, mop, dip, contaminate, water,	Vacuuming, floor cleaning
Professional cleaning		
7	pipe, piping, fuel_tank, pip, trough, circumferential, wiring_piping, diameter, worm_gear, minitube, tobacco_roll, circumferentially, internal_finned, crest, inner, tooth_stump, gauge, axial_direction, butt, lapse_predicted, pipes, levelling_bench, circumference, minitubes, society_for_biomolecular_screening, swedish_specification, fueling_pistol, supportive, readable, pitting,	Tank, tubes, pipes and sewers
218	water, tank, rock, bay, submerged, sediment, submerge, reverse_osmosis, purified_water, scavenge, ingredient, oht, pump, petroleum, chemical_mechanical_polish, underground, sink, borehole, purification, kitchen, buoyancy, rice, intrabay, liter, interbay, mpi_medium, cooking, immerse, purify, ballast,	Tank, tubes, pipes and sewers
223	housing, seal, sealing, bellow, inlet, outlet, opening, tight, leak, lip, sealed, interior, inner, evacuation, gland, venting, bellows, sealingly, diaphragm, evacuate, outer, enclose, annular, house, manifold, reactive_ion_etch_mode, rings, vent, inlet_outlet, closed,	Tank, tubes, pipes and sewers
Logistic systems		
144 ↑	shelf, automated, warehouse, inventory, automate, sort, mail, retail, shipping, fulfillment, facility, forklift, sorting, order, pallet, good, aisle, deadlock, retailer, logistics, management, tote, shipment, acceleration_slowdown, shelves, distribution, sale, depot, stock_keeping_unit, shelving,	Automated Guided Vehicles (AGVs) in manufacturing environments
Defense applications		
327 ↑	group, mission, game, weapon, surveillance, unmanned, lawn_mower, military, player, terrain, team, swarm, waist, soldier, helicopter, enemy, threat, reconnaissance, tactical, mobility, ugvs, squad, unmanned_air_vehicle, fly, sport, combat, ground, launch, drone, opponent,	Unmanned ground-based vehicles
Construction and demolition		
353	lay, fire, bake, explosive, containment, chill, supercritical, branch, laying, underground_pipeline, branch_pipe, detonation, detonate, baking, vernier, firing, rain_maker, wildfire, fire_fighter, buckle_arrestor, blasting_cap, photoresist, electric_discharge_machining, sewer, baked, disruption, sige, fire_fight, overlap_vernier, bake_chill,	Heavy/civil construction

Note: For each topic, its top 30 most relevant (i.e. with high likelihood and exclusivity, see [equation \(8\)](#) for formal definition) words are displayed. Up-arrows indicate that the diffusion of the respective topic follows a positive trend, vice versa for down-arrow. Topics within an area are sorted by their topic matching scores.

4. Results

4.1. Number of topics

As described in Section 3.3, we proceed with pre-processed texts by computing LDA models and plotting the corresponding scores on perplexity, coherence and exclusivity for different numbers of topics ranging from 10 to 500 with an interval of 10 (Fig. 5). The best model should have the lowest perplexity score (i.e. lowest prediction error), and the highest coherence and exclusivity. As one can see, the perplexity score improves with the number of topics, but the improvement is marginally decreasing. The same pattern is observed for the exclusivity, while coherence tends to aggravate for the growing number of topics. Choosing an optimal number of topics K that would maximize all three dimensions in such circumstances is not possible. An additional fourth criterion for choosing K is model complexity implying that with more topics it becomes increasingly difficult to overview and interpret them. For that reason, observing that after approx. 300 topics the three selection criteria start to change much slower, we chose K to be equal to 380 topics as for this number we observe a (small) local optimum both for perplexity and exclusivity. In addition, to demonstrate robustness of our results we henceforth report in parallel the topic model with 190 topics, i.e. exactly half of K . We chose the alternative value to be smaller to increase readability of the results, as charts for networks, diffusion curves and word clouds tend to be simpler for smaller K .

4.2. Matching topics to SR

Applying our matching procedure described in Section 3.4 and Appendix B on the constructed topics, we generate a distribution of topic matching scores ranging for the case of 380 topics between 0 (i.e. no candidate word overlapping between IFR reference texts on SR and our topics was identified) and 5.62. The largest topic matching score has been obtained for topic 377 matched with the agriculture IFR sub-area. We define the threshold value maximizing the total number of uniquely matched topics as 1.366 (see right plot in Fig. 6a). This is illustrated with the intersection of the black line showing the number of topics classified to only one SR sub-area and the red line demonstrating total number of topic matches to SR sub-areas including those cases where the same topic has been matched to two, three or more SR sub-areas. Taking a threshold value below 1.366 will result in topics simultaneously matched to more than one SR sub-area, while a threshold value above 1.366 would result in fewer uniquely matched topics having higher overlap with IFR SR descriptions. This threshold value resulted in 21 out of 380 topics matched to SR (see Table 3 for details on which topics have been matched to which SR sub-areas). Note that for 190 topics the topic matching scores vary between 0 and 5.59 (Fig. 6b). The largest matching score is again from the sub-area of agriculture for topic 105. The threshold value maximizing the number of uniquely matched topics is 1.280. This results in 20 matched topics (see Table C5 in Appendix C). Thus, the results of applying the matching procedure for 190 and 380 topics look very similar supporting the robustness of our results for different K .

The left panels of Figs. 7 and 8 show word clouds we matched to service robotics among the 380 and 190 topics, respectively (Figure C18 and C24).²³ For example, for model with 380 topics, topic 191 has words 'surgical' and 'surgeon' both very frequent and exclusive, while the words 'medical' and 'invasive' are neither as frequent nor as exclusive. The resulting diffusion curves can be found in Figure C19 and C25 Appendix C for the whole set of topics generated, while in right panels of Figs. 7 and 8 we show diffusion curves for SR topics among the 380 and 190 topics, respectively. Inspecting and comparing the

²³ Word clouds for the whole set of topics generated in both LDA models can be found in Appendix C.

Table 3

Metrics for networks of topics .

	1977- 2017	1981- 1985	1996- 2000	2011- 2015
LDA model with 380 topics				
Number of components	51	2	4	15
Clustering	0.41	0.26	0.32	0.36
Mean path length in LCC	95.46	4.45	8.55	29.92
Modularity	0.28	0.27	0.28	0.29
Degree centralization	0.31	0.35	0.28	0.35
Betweenness centralization	0.06	0.11	0.08	0.11
LDA model with 190 topics				
Number of components	22	3	7	16
Clustering	0.37	0.29	0.31	0.33
Mean path length in LCC	42.09	6.72	14.55	31.26
Modularity	0.33	0.34	0.35	0.36
Degree centralization	0.27	0.18	0.18	0.24
Betweenness centralization	0.13	0.10	0.08	0.14

generated word clouds and diffusion curves for the LDA models with 190 and 380 topics, two observations arise. First, the generated topics matched to SR in the two LDA models tend to strongly overlap. One example is the pair of topics with largest topic matching scores on milking robots (topics 377/380 and 105/190²⁴). Other good examples are topics on vacuuming and floor cleaning robots (284/380 and 6/190), robot assisted surgery (191/380 and 110/190) or tanks, tubes and pipes (7/380 and 134/190). Second, the topics matched to SR tend to be rising topics. Among 380 topics, their chances to be classified as rising are twice higher (38%, or 8 out of 21 topics) than for all topics constructed on average ($\approx 17\%$, or 63 out of 380 topics). For example, almost all of the topics matched to medicine (6, 191, 280, 308 and 311) have experienced a fast growth in the last ten years, while in the 1980s they were very small. There is also no falling topic among SR topics compared to 11% (41 out of 380) among all topics constructed. For 190 topics and twenty of them being matched to SR, the situation is similar with 40% (8 out of 20) vs. 20% (38 out of 190) for rising topics, and 5% (1 out of 20 matched SR) vs. 18% (34 out of all 190 topics) for falling topics.

4.3. Network perspective

The network generated for the whole period of 41 years (1977–2017) as well as three equally distant time windows of five years length are presented in Figs. 9 and 10.²⁵ As it was explained in Section 3.5, the density of our networks at any period by construction is 5%, while the absence of an edge between any pair of topics in our graph does not imply that those topics never appear together in our patent texts but that the frequency and prevalence with which they co-occur is not high enough to be considered as significant. The topics belonging simultaneously to 10% most central topics in terms of degree centrality (i.e. number of edges) and betweenness centrality (i.e. how often these topics are located on the shortest paths connecting any random pair of nodes in the graph) are marked with squared node borders. One can conclude that these topics represent some general technologies in the field of robotics appearing in many different patents and connecting different parts of the knowledge space in robotics together. In this sense, these can be compared to so called 'general purpose technologies' (see Korzinov and Savin (2018) for a recent discussion) that just like electricity or steam engine 'enable' development of many new inventions. Word

²⁴ Henceforth, for brevity reasons, we will indicate with ' $/K$ ' the LDA model the topic index belongs to.

²⁵ Since in the latter we take into account only particular sub-periods, presence of edges between topics can vary from period to period under consideration.

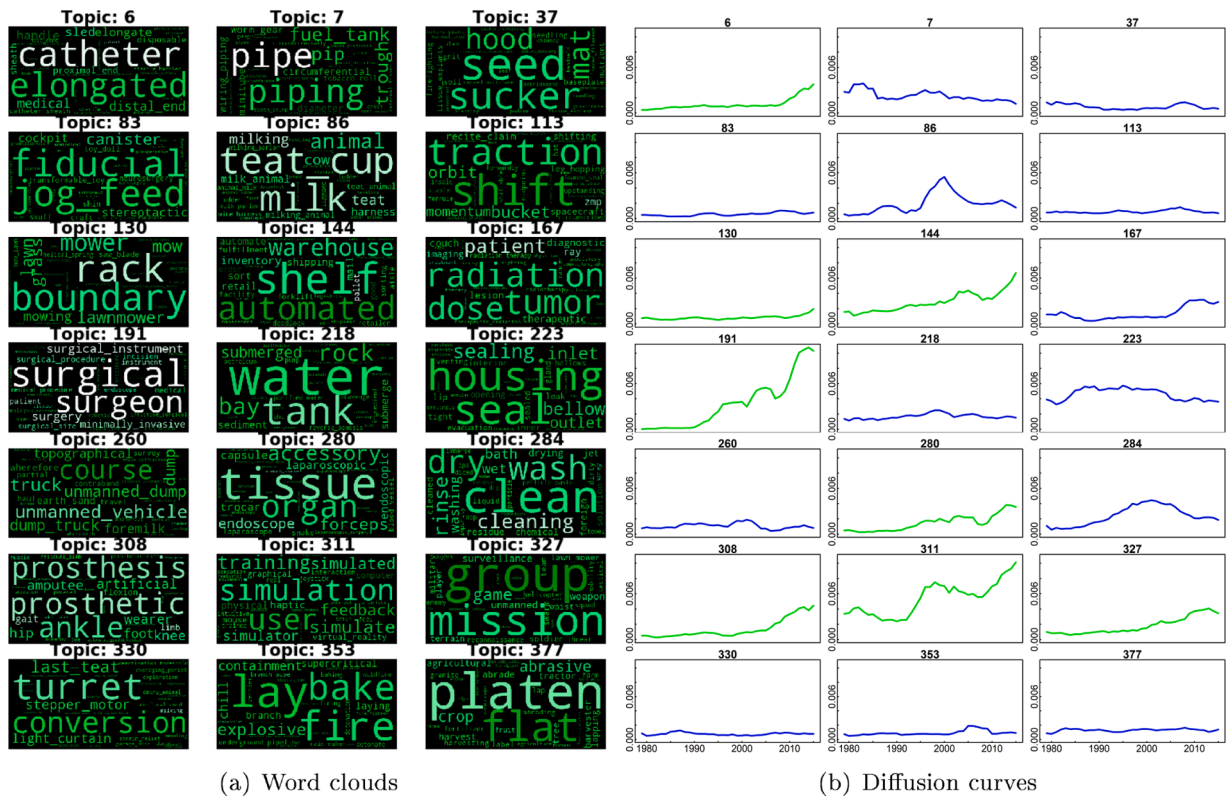


Fig. 7. 21 SR topics in the LDA model with 380 topics.

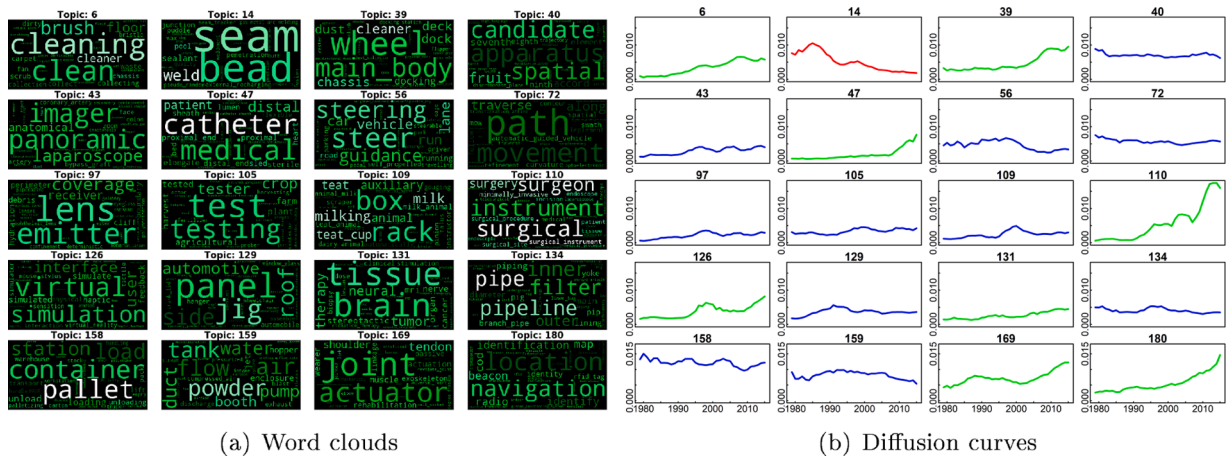


Fig. 8. 20 SR topics in the LDA model with 190 topics.

clouds of these topics are presented in Figs. C16 and C22 in Appendix C. Their visual inspection corroborates our intuition: these topics often deal with processing information, positioning and orientating in space, while words in the word clouds having dark font imply that they have a relatively low exclusivity.

In contrast, the isolated topics (or so-called 'isolates') represent parts of knowledge infrequently appearing together with other topics, having predominantly exclusive terms in their word clouds²⁶ and capturing what we call specialized application areas of robotic technologies. Take as an example topics 37 and 377 from the LDA model with 380 topics. Both are matched to SR (sub-area agriculture) and are isolates in the

network constructed for the whole period of 1977–2017. These topics are about technologies around seeding and harvesting (crops) using robots. These areas of application as such are quite specific and do not come often neither with general topics we mentioned earlier nor with other topics like those matched to SR or identified as isolates.

The majority of topics in our graphs are naturally in between these two extremes (most central topics and isolates). Identifying few edges for them we can see what other areas of robotic knowledge they are closest to, which can help in their interpretation. E.g. topic 191 (matched to SR sub-area robot-assisted surgery and therapy) in the LDA model for 380 topics is connected to topics 338 (containing many terms on telesurgery), 9 (addressing endoscopy and filtering) and 96 (manipulation and teleoperated machines). None of the latter three has been matched to SR via our matching approach. Thus, the network perspective gives a broader picture on what kind of technologies tend to be

²⁶ Word clouds of these topics are presented in Figs. C17 and C23 in the Appendix C.

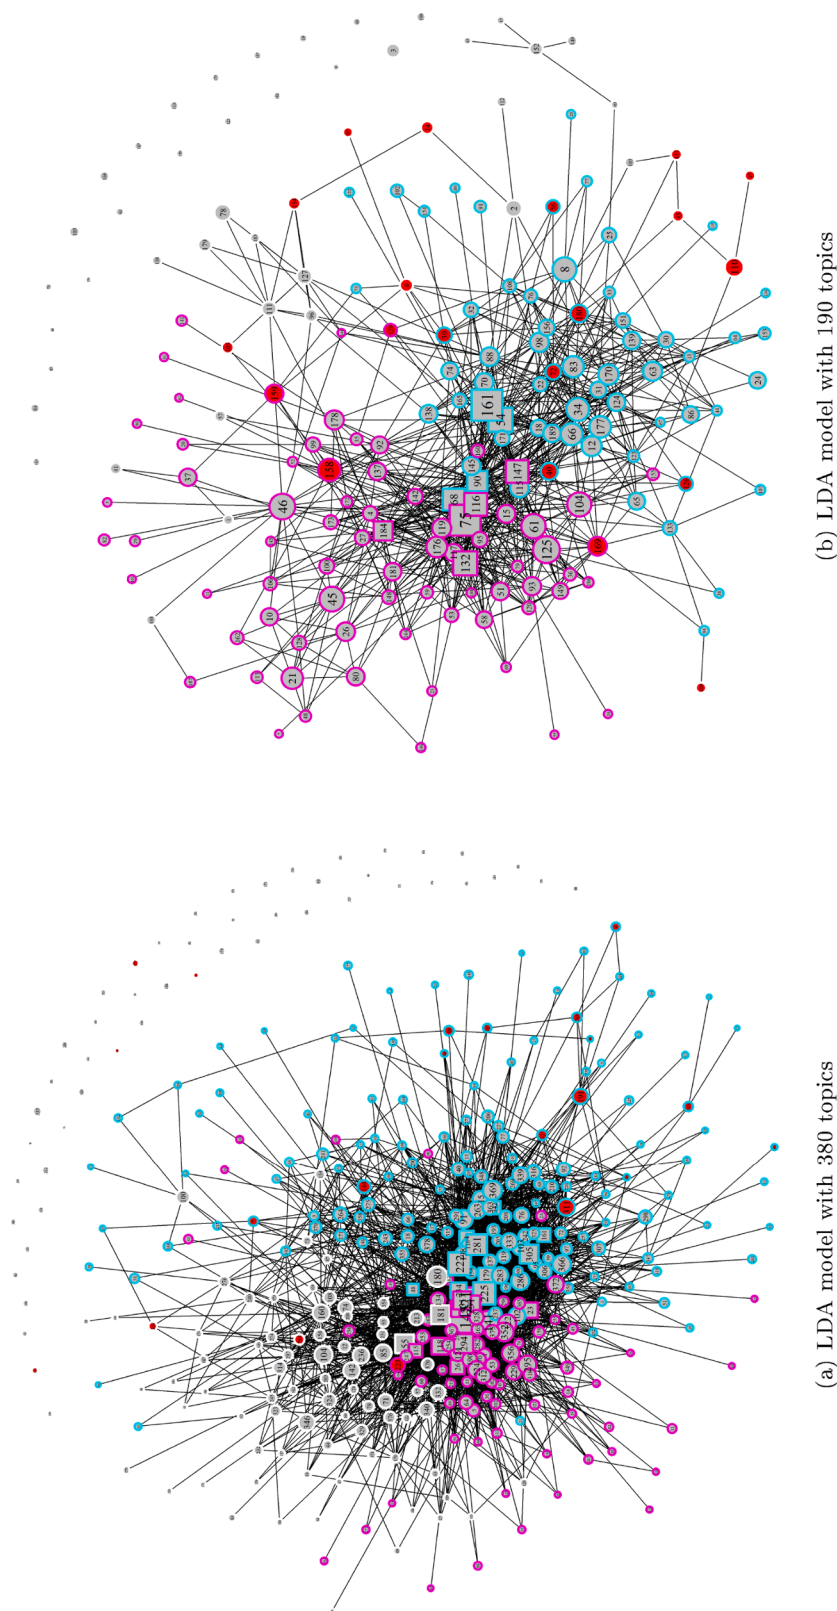


Fig. 9. Networks of topics for the period 1977–2017. *Note:* The size of the nodes represents the percent of the corpus of patents texts explained by the topic, while the presence of the edge captures the fact that the two topics tend to co-occur with large prevalence in many patent documents so that the chance to observe them together in a randomized null model is below 5% (see Section 3.5). Topics belonging to 10% most central in terms of degree and betweenness centrality are highlighted with squared frames, while topics matched to SR - with red color. Finally, cyan and magenta color are used to highlight the two clusters of topics on software and hardware, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

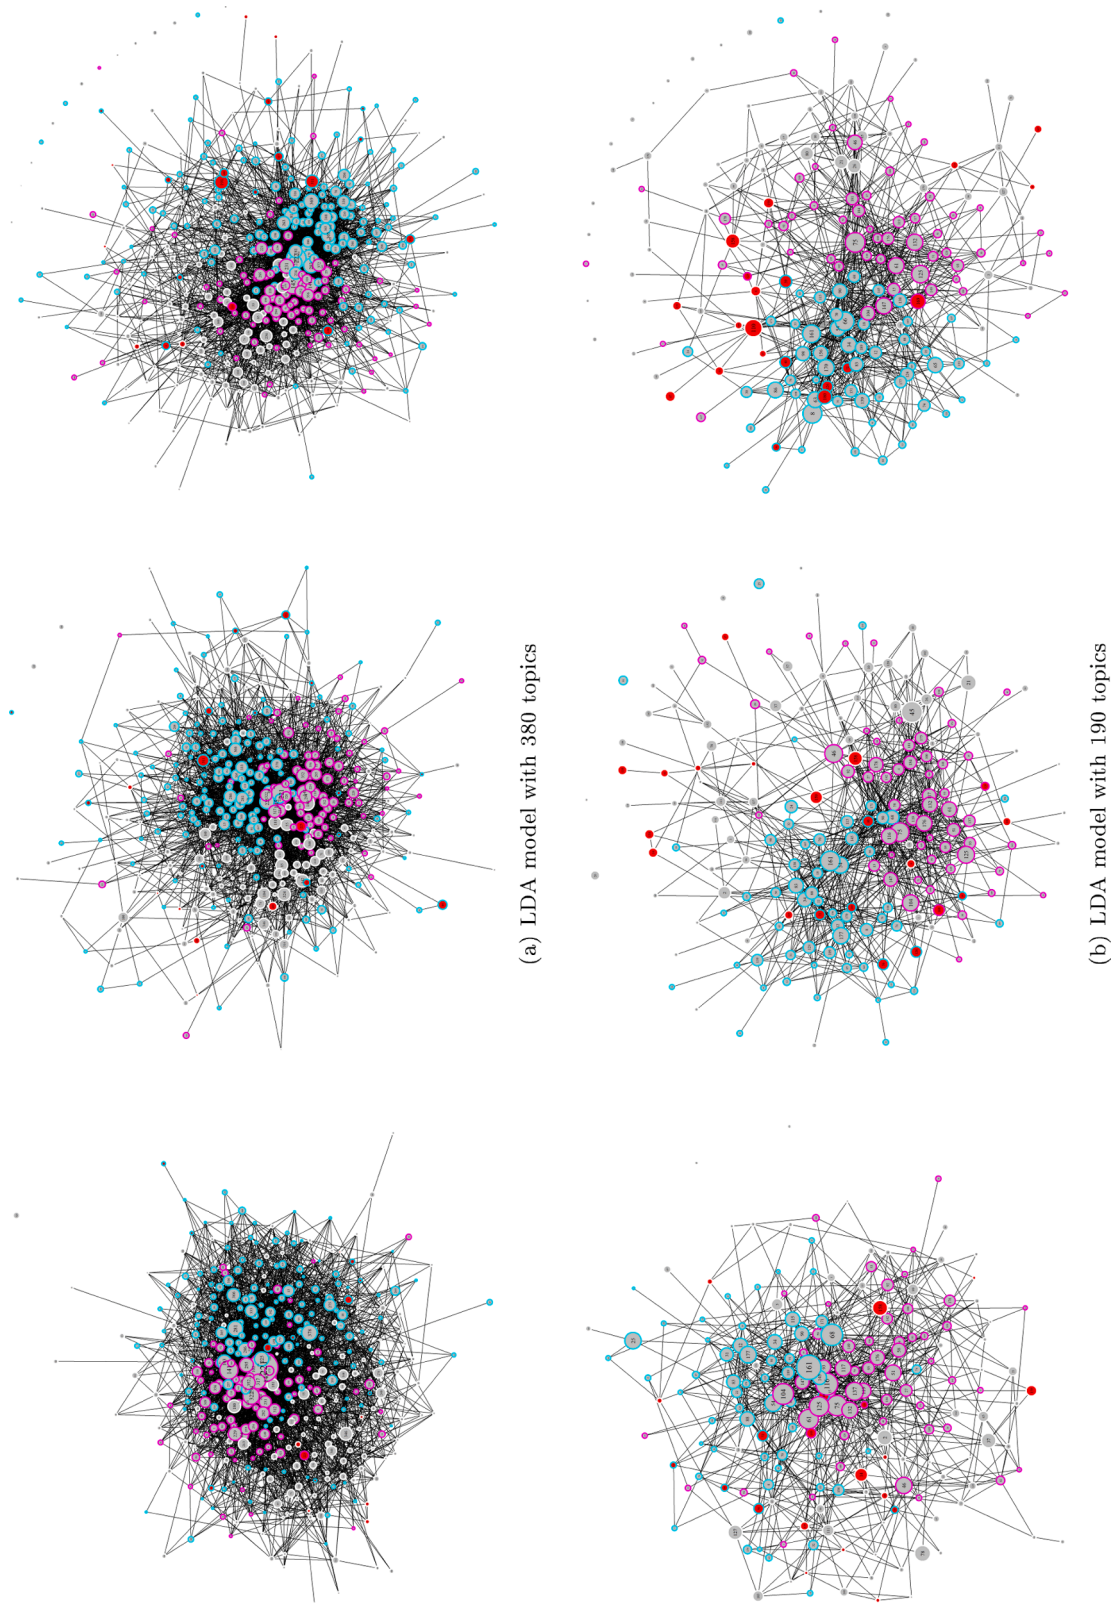


Fig. 10. Networks of topics for three time windows: 1981–1985, 1996–2000 and 2011–2015 *Note:* The size of the nodes represents the percent of the corpus of patents texts explained by the topic, while the presence of the edge captures the fact that the two topics tend to co-occur with large prevalence in many patent documents so that the chance to observe them together in a randomized null model is below 5% (see [Section 3.5](#)). Topics belonging to 10% most central in terms of degree and betweenness centrality are highlighted with squared frames, while topics matched to SR - with red color. Finally, cyan and magenta color are used to highlight the two clusters of topics on software and hardware, respectively. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. 11. Prevalence of topic clusters and SR over time. LDA model with 190 topics (left plot) and 380 topics (right plot).

recombined with each other in different patents to further extend the knowledge frontier.

Finally, to identify communities of nodes forming strong interconnected clusters in our graphs of topics, we apply the algorithm proposed by Newman (2006) based on leading eigenvector of the community matrix. This allows us to identify two large communities of nodes dominating the networks.²⁷ Looking closer on the word clouds of topics forming those two communities, one can quickly realize that these can be interpreted as clusters of technologies concentrated around either hardware technologies (e.g. body, movement, energy), or software technologies (e.g. receiving, processing and storing information). See word clouds with most central topics from each of the two cluster for both LDA models in Figure C14 and Figure C20 and C21 in Appendix C. We highlight these two clusters in Figs. 9 and 10 with different colors (magenta for hardware and cyan for software).

Interestingly, if we sum up prevalences of topics belonging to these two clusters and to SR over time, we clearly see on Fig. 11 that i) the software cluster overtook the hardware one in the share of patent documents, which confirms the rising importance of ICT; ii) service robotics as expected is a rapidly growing research field in the last few decades.²⁸

We calculate several graph indicators to quantify the generated networks and their development over time (Table 4). First, we see that the *number of components*, that captures how many nodes are isolated, increases over time. This we interpret as a sign that topics in robotics became more 'independent' and can form patentable knowledge with other topics. Second, The *clustering coefficient* of the network increases steadily. Apparently, the generated topics increasingly form densely connected cliques inside the largest connected component (LCC).²⁹ Third, the *mean path length* – measuring the average number of edges to connect any pair of randomly chosen topics in LCC (i.e. disregarding isolated topics) – clearly increased over time indicating a lack of short paths connecting the densely interconnected cliques. One can interpret

that as formation of hierarchy in the knowledge network, where topics from one area of knowledge tend to become isolated from the topics in the other areas. The intuition that our constructed networks are hierarchical is further supported by the *modularity* measure being around 0.3.³⁰ Concerning the importance of different topics in the network, we use the concept of network centralization that is concerned not with the overall connectedness but with the particular structure of relations and relative positions of the nodes. We estimate two commonly used centralization metrics. The *degree centralization* measures the concentration of linkages on few nodes (1 would be a star network and 0 – a complete graph). The *betweenness centralization* (again between 0 and 1) indicates the dependence on nodes that connect many other nodes. Both measures remain rather stable over time, at least if we look on the LDA model with 380 topics. Stability of these measures means that the asymmetry between topics in their centrality does not change and most central topics do not become more important in more recent years.³¹

5. Review, perspectives, and policy implications

The present paper extends the existing economic literature on robotic research in several directions. Using unstructured patent information on robotics that covers a time span of more than four decades and applying a topic modeling approach we discover latent topics in those texts, which we later match to SR sub-areas using a novel matching procedure. The latter is done by identifying words that appear both, in the lists of words associated to each topic with different probabilities and the descriptions of SR areas from IFR. Then we account for the frequency and exclusivity of these words (also compared to scientific literature outside IFR descriptions), and calculate a threshold level necessary to be reached to be identified as SR topic, which maximizes the number of unique SR topic matches. This way we contribute to the literature on topic identification and interpretation which is rather nascent when it comes to an automatic topic matching and labeling. Another key novelty of our analysis is the embedding of topics into a network perspective. This allows to capture interdependencies between robotic technologies by demonstrating which topics appear a lot together over many patent documents, and which do not. This allows to determine the relative position of SR topics within this network, observe

²⁷ In fact, the algorithm identifies more than two communities, but we focus only on the two largest ones in terms of number of nodes disregarding the smaller communities. This choice is primarily motivated by their size, while all other communities are much smaller and hard to classify to any area.

²⁸ One can see that in the LDA model with 190 topics the share of texts belonging to SR is considerably bigger than in the LDA model with 380 topics. This has to do with the fact that the model with less topics is less detailed in disaggregating patent texts into topics. While both measures are imperfect since they miss topics which overlap with IFR SR classification to a smaller extent than our matching threshold, we believe that the results for 380 topics shall be considered as a more accurate one.

²⁹ LCC comprises all the elements of the network, where any two nodes are connected to each other by a sequence of edges (path).

³⁰ This estimate is close to the one by Valverde et al. (2007) obtained for computed tomography of 0.41.

³¹ For 190 topics we observe an increase in both metrics of centralization in the last five-year window (2011–2015). Since the model with 380 topics is considered by us as more accurate one, we abstain from interpreting this result, as it may be an outcome of more rough classification of documents into smaller number of topics.

an increasing hierarchy of the network with high clustering and distinguish between topics that connect to many areas of knowledge enabling their recombination and topics that are rather isolated and focus on distinct application areas. Both these novelties are implemented in a data-driven way, independent of both experts' idiosyncratic knowledge and patents' structured information (such as concordances). We thus go beyond existing literature (e.g. [Kreuchauff and Korzinov \(2017\)](#)) that bases their SR patent identification strategy on expert judgment. Notice that patent classifications (such as International Patent Classification) also rely on expert knowledge and introduce new classes with a considerable time delay compared to first patents filed on that topic, thus being inherently weak to capture emerging trends in SR.

While examining our results, one should remember that they just as any other topic model are a result of stochastic optimization implying that running the LDA algorithm several times one can obtain every time (slightly) different outcome. Therefore, it is important to focus on the qualitative results of our analysis. To ensure their robustness, we compare two topic models: with 190 topics and 380 topics. Even though the two models have very different number of topics, the results they convey are very similar. First, they exhibit almost the same number of topics matched to SR (approx. 20). Second, predominant number of topics matched to SR are in field and medical robotics. Third, we find disproportionately large number of rising patterns among the diffusion curves of topics matched to SR. Fourth, very similar structural properties of the networks of topics (two large communities reflecting hardware and software, position of topics matched to SR in the periphery of the network, similar level and dynamics for network metrics such as clustering, modularity and centralization). Fifth, rising trend in the share of service topics and software cluster among topics on robotics. These results convince us that the patterns we observe for service robotics are robust to alternative specification of LDA parameters.

Our approach enables us to replicate some stylized facts of robot technology transformations as identified by the use of structured patent data and may thus be compared to analysis based on well established concordances ([WIPO \(2008\)](#) for patent-technology mapping, [Eurostat \(2014\)](#) for patent-industry mapping). Among these facts are the relative shift of focus in robotics technology from hardware to software, the increasing importance of robotics in medicine (see [Table A12](#) in the Appendix), agriculture and logistics (see discussion in [Section 2.1](#)). With the proposed method we can look closer into specific SR sub-areas. In case of medical robotics we see the increasing importance of surgery and therapy but also - and this is not yet common knowledge - a strong trend of robot technologies related to rehabilitation.³² Besides, even specific areas of application can be identified and spotted. For example, we find the term 'laparoscopy',³³ which is an operation performed in the abdomen or pelvis using small incisions (usually 0.5-1.5 cm) with the aid of a camera. In the technology description of the SR report, the word 'laparoscopic' shows up while describing various types of operations carried out by robots ([IFR, 2018b, 136](#)). Associated topic 280/380 has been matched to medical robots (see [Table 3](#)).

We are able to interpret results of our analysis in the light of economic theories and recent technological trends. Most importantly, we demonstrate automation potential materializing within services, an economically important sector which until today has not yet seen much automation. While robot adoption in manufacturing industries mostly has the advantage to relieve humans from dirty, dull and dangerous tasks, in services non-standardized human-to-human interaction processes are essential. The use of SR may help to overcome labor shortages

in e.g. elderly care, logistics and cleaning, which is particularly important in the light of recent demographic trends in industrialized countries.

Related to robot markets, we can link the dynamics of SR patenting to SR sales. Comparing sales data and diffusion curves results in a robust finding that SR in medicine and logistics are the leading sectors of service robotics.

The paper at hand has several policy implications. Looking in the past, we see that robot patenting activity has been less susceptible to the financial crisis (2007–2009) than robot investment. However, in the light of the current COVID-19 pandemic the potential to reduce contagion risks might not only spur robot research but also become a driver of robot demand.³⁴ It is natural to assume that this will foster sales of service robots also in the long term and spur private investment in robotic technologies. Governments could support this process by complementary basic research funding and subsidizing applied research in areas where market incentives are not strong enough to fulfill actual needs.

For example, related to care, lack of qualified labor is evident. Increasing the automation intensity here would help to solve this problem. An intensified use of robots would furthermore alter the associated work environment as well as the skill requirements of caretakers making the job attractive for technology-oriented workers. As a consequence, the educational system has to be adjusted to the new needs and opportunities. To cope with this process in a timely and efficient manner, a sound understanding of SR is important. And this is where our study comes into play as we show that information out of patents can be used as early signal prior to the actual transformation. Furthermore, in light of the recent progress in NLP and ML methods and growing amount of data being produced in sensitive application areas like healthcare, adequate conditions related to data privacy and data security have to be established. The governments should especially think about how high quality data may be made available for independent scientists to avoid monopolization of this resource and thus to create a level playing field at the global level.

6. Conclusions and outlook

There is a vast amount of studies on robotics in the economic literature. However, most of them are related to IR with a focus on labor market implications. Services, in contrast, are still rarely addressed although here the automation potential is huge. This gap in the literature is mostly due to a lack of reliable data and the difficulties to quantify emerging technological fields. We shed new light here by applying topic modeling and graph theory to textual patent data. Furthermore, given that topic labeling and interpretation is a challenging task lacking objective data-driven methods, we proposed and applied a novel approach that uses external reference texts on SR robotics, namely technology descriptions of SR from the IFR.

Our results reaffirm that at present various disciplines are involved in technological development of robotics. Among the most prominent are computer science, automation control, electrical and mechanical engineering and bio-medicine ([Analytics \(2018\)](#)). Today's robots are increasingly characterized by more intelligent components (e.g. smart grippers), greater connectivity (e.g. plug & play interfaces, cloud computing) and are easier to use (e.g. programming by demonstration). Although these technologies are also applied in IR, the future of modern robots happens increasingly beyond predefined environments. As a consequence, the areas of application also change constantly over time: logistics, medicine, agriculture, construction, cleaning. The rapid technology advances will likely make robots applicable also in other sectors that until now have not yet seen much automation. These technologies might help to overcome a phenomenon which has become known as

³² Within medical applications we are even able to further differentiate the usage of robots, e.g. the increasing role of rehabilitation systems (topic 308/380) and surgery (topic 311/380 focuses more on the hardware perspective while topic 191/380 relates to the complementary software).

³³ From ancient Greek *λαπαρά* (lapara) 'flank, side', and *σκοπεω* (skopeo) 'to see'.

³⁴ [Caselli et al. \(2020\)](#) have recently shown that robotization indeed facilitates social distancing and lowers the risk of contagion.

'Baumol's cost disease'³⁵ or the phenomenon of secular stagnation, i.e. higher productivity in service sectors which is due to automation contributing to aggregate productivity growth.

This study has a few limitations, which can also be seen as directions for further research. First, we could not find comparable sources of descriptions to automatically label the remaining (non-SR) topics. Applying our topic matching on a broader classification of robotics is therefore welcome. Second, patent data covers mainly technical side of technologies in robotics. One could go for a broader perspective and use LDA for other related data sources such as scientific publications, funding programs, newspapers and other media coverage. Such data would address societal discourses and political debates about robotics to elicit not technologies but cultural differences, such as arguments in favor and against robots as drivers or doctors. This way one could draw a comprehensive picture of SR acceptance including the perspectives of all stakeholders within the technology system. Third, we did not look on the

position of countries in mastering those technological topics. It is worthwhile to assess revealed technological advantages of each country in different SR sub-areas and how these were changing over time. This would be of particular interest for policy makers to design their supporting measures for local companies competing within the international technology markets. Furthermore, since topic diffusion curves may be linked to sales dynamics, one could extend the analysis by building market forecasts based on the patenting activities.

CRediT authorship contribution statement

Ivan Savin: Conceptualization, Formal analysis, Methodology, Writing – original draft. **Ingrid Ott:** Conceptualization, Formal analysis, Writing – original draft. **Chris Konop:** Conceptualization, Formal analysis, Methodology.

Appendix A. Classification of robots: ISO standard and IFR operationalization

A1. Robot categories according to standard ISO 8373:2012(en)

(ISO 8373:2012, 2.6) A robot is an actuated mechanism programmable in two or more axes with a degree of autonomy, moving within its environment, to perform intended tasks. Autonomy in this context means the ability to perform intended tasks based on current state and sensing, without human intervention.

(ISO 8373:2012, 2.10) A service robot is a robot that performs useful tasks for humans or equipment excluding industrial automation application. Note: The classification of a robot into industrial robot or service robot is done according to its intended application.

(ISO 8373:2012, 2.11) A personal service robot or a service robot for personal use is a service robot used for a non-commercial task, usually by lay persons. Examples are domestic servant robot, automated wheelchair, and personal mobility assist robot.

(ISO 8373:2012, 2.12) A professional service robot or a service robot for professional use is a service robot used for a commercial task, usually operated by a properly trained operator. Examples are cleaning robot for public places, delivery robot in offices or hospitals, fire-fighting robot, rehabilitation robot and surgery robot in hospitals. In this context, an operator is a person designated to start, monitor and stop the intended operation of a robot or a robot system.

(ISO 8373:2012, 2.15) A robot system is a system comprising robot(s), end-effector(s) and any machinery, equipment, devices, or sensors supporting the robot performing its task.

NACE code	NACE code description	1981-1985	1996-2000	2011-2015
28.4	Manufacture of Metal Forming Machinery and Machine Tools	0,4394	0,3011	0,2374
26.5	Manufacture of Instruments and Appliances for Measuring, Testing and Navigation; Watches and Clocks	0,1248	0,0878	0,0800
26.51	Manufacture of Instruments and Appliances for Measuring, Testing and Navigation	0,0966	0,0619	0,0937
28.29	Manufacture of Other General-Purpose Machinery N.E.C.	0,0841	0,0647	0,0529
26.1	Manufacture of Electronic Components and Boards	0,0502	0,1531	0,0579
28.1	Manufacture of General-Purpose Machinery	0,0304	0,0467	0,0560
28.9	Manufacture of Other Special-Purpose Machinery	0,0281	0,0427	0,0408
32	Other Manufacturing	0,0209	0,0171	0,0143
29.1	Manufacture of Motor Vehicles	0,0154	0,0204	0,0243
32.5	Manufacture of medical and dental instruments and supplies	0,0119	0,0253	0,0741
28.23	Manufacture of Office Machinery and Equipment (Except Computers and Peripheral Equipment)	0,0105	0,0288	0,0284
26.2	Manufacture of computers and peripheral equipment	0,0099	0,0224	0,0985
25.3	Manufacture of Steam Generators, Except Central Heating Hot Water Boilers	0,0093	0,0053	0,0019
25.2	Manufacture of Tanks, Reservoirs and Containers of Metal	0,0072	0,0024	0,0011
32.9	Manufacturing N.E.C.	0,0048	0,0088	0,0066
27.1	Manufacture of Electric Motors, Generators, Transformers and Electricity Distribution and Control Apparatus	0,0047	0,0058	0,0093
27.9	Manufacture of other electrical equipment	0,0044	0,0032	0,0027
27.33	Manufacture of Wiring Devices	0,0036	0,0089	0,0029
28.22	Manufacture of Lifting and Handling Equipment	0,0035	0,0026	0,0047
26.7	Manufacture of Optical Instruments and Photographic Equipment	0,0035	0,0047	0,0058

Fig. A12. Robot patent–industry assignments and their relative importance across three time windows Note: Results are provided in weighted patent counts mapped according to the PATSTAT database. Color highlighting only refers to those industries that we address throughout our study. Green (orange) color highlighting indicates an increase (decrease) in relative importance over time, which can be seen e.g. in computers and peripheral equipment or medical and dental instruments (e.g. machine tools).

³⁵ Baumol and Bowen (1966) observed that rapid productivity growth in some sectors (e.g. in manufacturing) relative to other sectors (e.g. service industries) could result in a 'cost disease' at the aggregate level if the slowly growing sectors constitute a large part of the economy. Baumol (2012) provides a recent discussion with a key focus on healthcare, while Aghion et al. (2017) adapt this reasoning to AI and automation.

A2. Robot patent-industry assignments based on PATSTAT mapping

Appendix B. Further details on the methods used

B1. Data pre-processing

The following stopwords have been added to the NLTK's list of English stopwords: 'would', 'could', 'figure', 'whether', 'wherein', 'respectively', 'primary', 'certain', 'etc', 'described', 'herein', 'first', 'second', 'third', 'fourth', 'fifth', 'moreover', 'furthermore', 'according', 'yes', 'present', 'disclosure', 'preferably', 'feature', 'like', 'embodiment', 'comprising', 'preferred', 'example', 'exhibit', 'develop', 'demonstrate', 'comprise', 'consist', 'prefer', 'prefering', 'preferred', 'otherwise', 'alternative', 'alternatively', 'discuss', 'generally', 'exemplary', 'substantially', 'typically', 'additional', 'upon', 'suitable', 'feature', 'aspect', 'design', 'employ', 'create', 'multiple', 'understood', 'non', 'depict', 'previously', 'partially', 'say', 'take place', 'great', 'suitable', 'design', 'generally', 'exemplary', 'aspect', 'combination', 'accordance', 'appreciate', 'arrangement', 'disclose', 'typically', 'understood', 'multiple', 'skilled', 'configuration', 'create', 'employ', 'upon', 'direct', 'mechanic', 'alter'.

The normalized pointwise mutual information (NMPI) score is calculated as follows:

$$\text{score}(\text{word}_a, \text{word}_b) = \frac{\log(\text{Pr}(\text{word}_a, \text{word}_b) / (\text{Pr}(\text{word}_a) \times \text{Pr}(\text{word}_b)))}{-\log(\text{Pr}(\text{word}_a, \text{word}_b))} \quad (1)$$

This score sets the probability of two words occurring together in relation to the probability of them occurring together in case of independence. A score of -1 (in the limit) means two words are never occurring together, 0 in case of independence (they occur together as often as expected, based on their independent probability), and +1 in complete co-occurrence. In our case, two words will be merged into a bi-gram with a "_" symbol, if the score is above 0.5.

B2. Formal description of LDA and the parameters used

More formally LDA can be described as on Fig. B13. The shaded circle represents the observed data ($w_{(d,n)}$, n th word in each document d). The unshaded circles denote latent (unobservable) variables: $z_{(d,n)}$ - the topic assignment for the n th word in document d (or, alternatively, the assignment of words to topics), θ_d - the topic proportions for the d th document, and $\psi_{(1:K)}$ - topic distributions over the vocabulary. Arrows indicate the conditional dependencies between variables, while frames (the boxes in the figure) refer to repetitions of sampling steps, with the variable in the lower right corner referring to the number of samplings: number of documents D and the number of topics K . Thus, the inner frame over $z_{(d,n)}$ and $w_{(d,n)}$ represent the repeated sampling of topics and words until N words have been generated for each document d ; the frame surrounding θ_d illustrates the sampling of a distribution over topics for each document d for a total of D documents; the frame surrounding ψ_k illustrates the repeated sampling of word distributions for each topic assignment until the word probabilities of K topics have been generated.

We have used the parallelized Latent Dirichlet Allocation from gensim, an open-source library for unsupervised topic modeling and natural language processing, using modern statistical machine learning implemented in Python. To estimate the LDA model, several simplifying assumptions have to be made. First, θ is assumed to be a random draw from a uniform Dirichlet distribution with scaling parameter α . Second, ψ is assumed to be a random draw from a Dirichlet distribution with parameter β . Following Griffith and Steyvers (2004) and the default settings of the LDA gensim package (<https://radimrehurek.com/gensim/models/ldamodel.html>), we set $\alpha = 1/K$, while β as a prior is learned by LDA from the data. For $T > 100$ $\beta \ll 0.01$, which shall result in a fine-grained decomposition of the texts into topics that address specific areas of knowledge.³⁶ Then, fixing the number of topics K ex ante and using Gibbs sampling one can iteratively approximate the hidden LDA model using the Markov Chain Monte Carlo approach converging to the target distribution by repeated sampling from it (Griffith and Steyvers, 2004).

Note also that the number of documents to be used in each training chunk of the LDA algorithm (so-called 'chunksize') was set to 1800, the number of passes through the corpus during training ('passes') equal to 20 and the maximum number of iterations through the corpus when inferring the topic distribution of a corpus ('iterations') equal to 1000. All other parameters were set to default values.




Fig. B13. Graphical model representation of LDA (based on Blei et al. (2003)).

³⁶ In our sensitivity analysis we experimented with setting either both these parameters or just one of them to higher values that are popular in the literature, such as $\alpha = 0.1$ or $\beta = 0.1$ or $\beta = 0.01$ (Chen et al., 2017; Huang et al., 2017; Kaplan and Vakili, 2015). However, those alternatives resulted in worse performance of topics in terms of perplexity and exclusivity (discussed later in this paper), and so we stick to the default setting.

B3. Determining the number of topics

The *perplexity score* measures the ability of an LDA model estimated on a subset of documents (training data) to predict the word distribution in the remaining subset of documents (testing data). It is defined as the exponential of the negative normalized predictive likelihood under the model. A lower perplexity score indicates that the model has better generalization performance. Formally, for a testing data (D_{test}) with T_1 documents, the perplexity score is equal to

$$\text{perplexity score}(D_{test}) = e^{-\frac{\sum_{d=1}^{T_1} \log p(w_d)}{\sum_{d=1}^{T_1} N_d}}, \quad (2)$$

where N_d is the number of words in document d , w_d is a vector of all the words in document d , and $p(w_d)$ is the probability of observing the word vector w_d in document d given the LDA model estimated from the training data.

Semantic coherence is a criterion developed by (Mimno et al., 2011) and is maximized when the most probable words in a given topic frequently co-occur together. The authors show that the metric correlates with human judgment of topic quality. Formally, let $D(v, v')$ be the number of times that words v and v' appear together in the same document. Then for a list of the M most probable words in topic k , the semantic coherence for topic k is given as:

$$C_k = \sum_{i=2}^M \sum_{j=1}^{i-1} \log \left(\frac{D(v_i, v_j) + 1}{D(v_j)} \right) \quad (3)$$

A smoothing count of 1 is included to avoid taking the logarithm of zero.

Finally, drawing on previous work on exclusivity and diversity in topic models (Bischof and Airolidi, 2012), the *exclusivity* of a topic measures if words with high probability under topic i have low probabilities under other topics. If so, we can conclude that topic i is exclusive. A topic that is both cohesive and exclusive is more likely to be semantically useful. More concretely, for topics $T_n \in T$ and words $w_n \in T_i$, we measure the exclusivity by:

$$\text{exclusivity} = |T| \times \sum_{T_1 \in T} \sum_{w_1 \in T_1} \sum_{T_2 \in T} \sum_{w_2 \in T_2 \cap T_1} \frac{1}{\text{weight}(w_1) \times \text{weight}(w_2)} \quad (4)$$

B4. Plotting diffusion curves for LDA topics

To produce diffusion curves, we follow Lenz and Winker (2020) in quantifying the probability that a topic T_i appears in the corpus of patent texts C for a certain year t as:

$$P(T_i, C_t) = \frac{\sum_{d \in C_t} P(T_i, d)}{D_t}, \quad (5)$$

where d stands for patent documents in which we sum the prevalences of that topic in that time period, normalized by the overall number of documents in period t , D_t . Hence, at any period of time, the probabilities of all topics sum up to one. In addition, we smooth the diffusion curve by applying a five-year moving average window so that at any period the equation (5) is estimated not just for the year t , but also for two years preceding t and two years following it.

Furthermore, to ease visual inspection of the diffusion curves, we classify the lines into three categories. To be classified as a 'rising' topic (with green color), the following inequality should hold:

$$\sum \Delta P(T_i, C_t) = \sum (P(T_i, C_t) - P(T_i, C_{t-1})) \nu \times \sum |\Delta P(T_i, C_t)| \quad (6)$$

where $\nu \in [0, 1]$. In other words, over the period of consideration the topic should predominantly rise and not fall. Reversely, to be classified as 'falling' topic (with red color), the inequality in (6) must hold with the opposite sign. If none of the two conditions is fulfilled, the topic is classified as 'in-between' (with blue color).³⁷

B5. Topic matching

Technically, the matching procedure consists of the following steps:

1. Define the reference texts: take IFR report on SR (IFR, 2018b). Lemmatize these texts and replace capital letters with small letters. This results in reference texts covering 49 sub-areas (compare Table 1).
2. Identify 'candidate words': search for words which appear simultaneously among 1000³⁸ frequent words for each topic and in the SR reference texts.
3. For each candidate word w_j of topic t and a reference text q , determine a word matching score according to (7):

³⁷ The value of ν is chosen to be equal to 0.3 to separate the diffusion curves into the three groups to be as distinct as possible. As long as ν is positive and not very close to zero (to avoid classifying as rising those topics where the sum of ups and downs is approximately the same) it does not affect our later conclusions.

³⁸ We look on the 1000 words having highest probability to be assigned to each individual topic. Looking beyond 1000 most likely words does not improve the results as the probabilities of the remaining words are virtually equal to zero.

$$\text{word matching score}(w_j, t, q) = \frac{\text{relevance}(w_j|t)}{c^{\theta-1} \times p(w_j)} \times \sum_{i=1}^n \frac{1}{i} \quad (7)$$

where

- $\text{relevance}(w_j|t)$ is the relevance of the word to the topic according to (8) Sievert and Shirley (2014), which captures how exclusive this word is in this topic.

Let $p(w_j|t)$ be the probability of word j in topic t as resulted by LDA and $p(w_j)$ the marginal probability of word j in the empirical distribution.

The lift is defined as the ratio of a words' probability in a topic to its marginal probability across the corpus, $\frac{p(w_j|t)}{p(w_j)}$ (Taddy, 2012). Then the relevance of word j for topic t is defined as

$$\text{relevance}(w_j|t) = \lambda \times \log(p(w_j|t)) + (1 - \lambda) \times \log\left(\frac{p(w_j|t)}{p(w_j)}\right), \quad \lambda \in [0, 1] \quad (8)$$

We chose λ equal 0.5, meaning equal weighting of the word's probability for topic and its lift. The idea is to reduce the weight of words appearing in other topics.

- $\theta \in [1, 49]$ is the total number of SR sub-areas the word occurs in. This way we discriminate words appearing in multiple SR sub-areas applying a nonlinear penalty $c^{\theta-1}$. We experimented with several scalars for c above one. Larger values of c apply a stronger penalty on words that appear in multiple SR sub-areas reference texts, thus radically reducing the weight of these words and thus shifting the focus to words which are more unique for particular SR sub-areas. We decided for $c = 4$ as increasing it further was not changing the results. as
- ρ is a generality score of the word based on its frequency in the English language and derived from the word frequency database composed by Mark Davies, which describes how often a word appears in the English language per million words. We normalize this frequency by our corpus' length and take its log to derive ρ for each word we matched.³⁹

This way we take into account that the reference texts represent only a fraction of all technologies (i.e. not covering IR and technologies outside the field robotics) and do not sufficiently penalize common words being matched by our algorithm (e.g. 'cup').

- n is the number of times the word appears in the same SR sub-area. Thus, every time a word appears in the same reference text, its marginal score will be increased with a declining margin based on the frequency it appeared.⁴⁰

Thus, the word matching scores in equation (7) are unique for each word W_j and each topic – SR sub-area pair.

4. Determine the topic matching score of a topic t belonging to a reference text q by summing the word matching scores of all relevant candidate words appearing in topic t and a reference text q :

$$\text{topic matching score}(t, q) = \sum_{w_j \in t \cap q} \text{word matching score}(w_j, t, q) \quad (9)$$

5. Derive a threshold for the resulting topic matching scores. To this end, consider the distribution of values for each topic and each SR sub-area and choose one that maximizes the number of topics classified to at most one SR sub-area. The threshold value is calculated on the interval between zero and $\max(\text{topic matching score}(t, q))$, where smaller values are associated with many topics matched to several SR sub-areas simultaneously (i.e. not uniquely), while increasing the threshold value reduces the amount of multiple matches until eventually each matched topic may be unequivocally be assigned to a single SR sub-area. This way, the threshold value is dependent on the distribution of topic matching scores: the higher the values of these matches indicating the overlap with multiple SR sub-areas, the higher the threshold value. We prefer to use this rule to define the threshold value as it best suits our purpose to identify as many topics capturing technologies from SR as possible, while remaining data-driven. Fig. 6 illustrates the derivation of the threshold value.

B6. Network space of topics

Remember that one of the outputs from applying LDA is a matrix of topic prevalences across the documents. Thus, each topic is a vector of weights, bounded between 0 and 1, of length D (number of our patents being 22'927), so that cosine similarity of topics a and b is:

$$\text{Cosine similarity}_{ab} = \frac{ab}{\|a\| \|b\|} = \frac{\sum_{i=1}^D a_i b_i}{\sqrt{\sum_{i=1}^D a_i^2} \sqrt{\sum_{i=1}^D b_i^2}} \quad (10)$$

Appendix C. Further results

C1. LDA model with 380 topics

³⁹ <https://www.wordfrequency.info/>. The list of around 100'000 words and their forms together with the frequency of use in the academic literature of the corpus of contemporary American English.

⁴⁰ Thus, the first time the word appears, its score is one, whereas the n -th time the word appears, its score is $\frac{1}{n}$.



Fig. C14. Word clouds of 16 most central topics belonging to hardware cluster among 380 topics. *Note:* Topics are listed in the order of centrality. Topic 223 is SR topic.

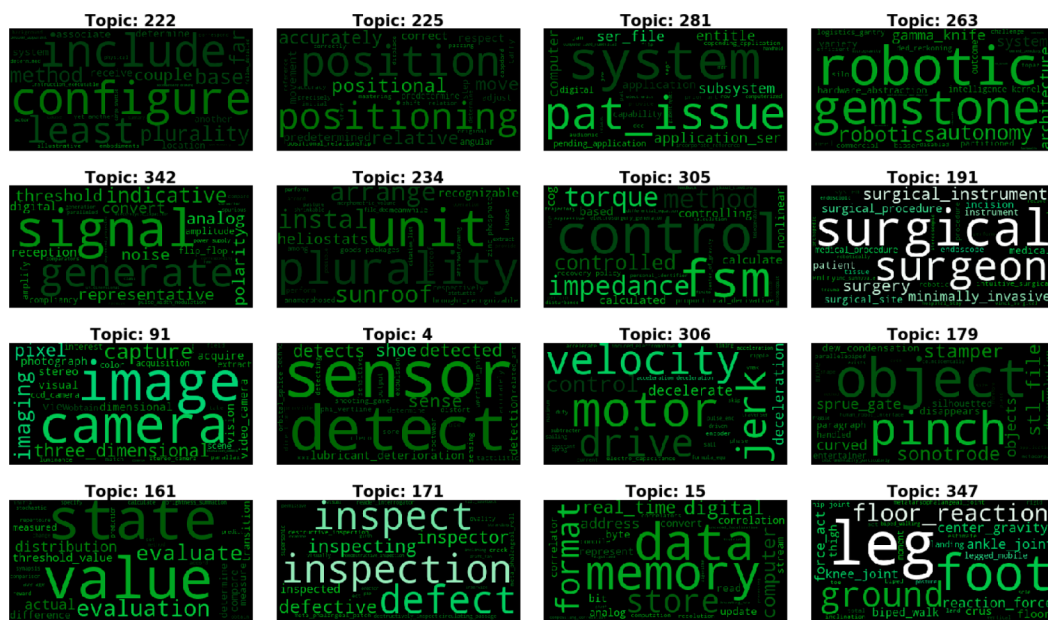


Fig. C15. Word clouds of 16 most central topics belonging to software cluster among 380 topics *Note:* Topics are listed in the order of centrality. Topic 191 is SR topic.



Fig. C16. Word clouds of 21 most central topics among 380 topics Note: Topics are listed in the order of centrality.



Fig. C17. Word clouds of isolated topics among 380 topics Note: Topics 37, 83, 353 and 377 belong to SR.




Fig. C18. Word clouds of 380 topics.




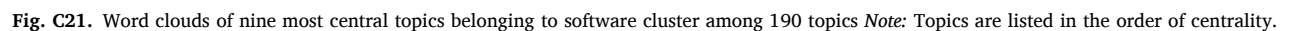
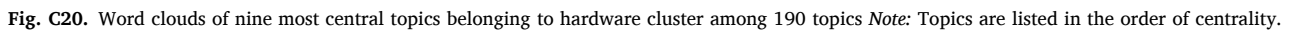
Fig. C19. Diffusion curves for 380 topics .

C2. LDA model with 190 topics

Table 4

Topics from LDA model with $K = 190$ matched to SR sub-areas. For each topic, its top 30 most relevant (i.e. with high likelihood and exclusivity, see Eq. 8 for formal definition) words are displayed. Up-arrows indicate that the diffusion of the respective topic follows a positive trend, vice versa for down-arrow (which only apply for matches within the 190 topic model). Topics within an area are sorted by their topic matching scores.

Topic	Topic terms (30 most relevant)	Matched SR sub-area
Field Robotics		
105	test, testing, tester, crop, agricultural, plant, harvest, tested, farm, harvesting, contactor, fertilizer, prober, actor, nondestructive_testing, pest, retinal, nutrient, soil, retinal_ganglion, destructive_testing, stimulus, pesticide, subsurface, farmer, corn, flower, probers, tractor, aerial_farm,	Agriculture
109	box, rack, milking, milk, auxiliary, teat_cup, animal, teat, milk_animal, teat_animal, animal_milk, pinion, scraper, dairy_animal, automatically, talus, gouging, instructor, microbial_agent, esf, mtp, relevant, compartment, console, subtalar, implement, head_mount_display, rear, watch, farmer,	Milking robots and livestock robotics
56	steer, steering, guidance, vehicle, lane, car, run, running, road, self_propelled, parking, driver, travelling, steerable, park, cage, pedal, weighted_constant, race, wheel, driverless, reconfigurable_toy, droplet_ejection, passenger, amusement, propulsion, lateral, ahead, appraisal, brake_pedal,	Agriculture
40	apparatus, candidate, spatial, fruit, accord, seventh, eighth, element, ninth, discriminant_plane, trajectory, region, predicted, sign_language, sixth, tenth, eleventh, condition_satisfied, twelfth, move, recognize, satisfy, thirteenth, condition, possibility, shade, action, predict, toner, tree,	Agriculture
Inspection and maintenance systems		
134	pipe, pipeline, inner, filter, outer, piping, main, branch_pipe, pip, diameter, yoke, lining, pig, liner_bag, smoking_article, circumferential, lay, leak, oil, underground, tubular, inside, rig, wall, interferogram, liner, dig, laying, interferometry, flux_leakage,	Tank, tubes, pipes and sewers
14 ↓	seam, bead, weld, sealant, junction, pool, penetration, puddle, external_recharging, wax, seam_tracker, fusion, pseudo_random, quality, weldment, gas_metal_arc_welding, butt, profiler, irregularity, corrugate, join, geometry, lateral, molten, slag, orbital_welder, gtaw, width, toolspeed, conformation,	Tank, tubes, pipes and sewers
159	air, powder, tank, flow, water, duct, pump, booth, enclosure, hopper, jet, discharge, exhaust, compressed_air, blast, pressurized, intake, stream, decontamination, orifice, spray_booth, airflow, flow_rate, fan, blower, curtain, blasting, interior, cyclone, mist,	Tank, tubes, pipes and sewers
Construction and demolition		
129	panel, jig, side, roof, automotive, automobile, hanger, window_glass, assemble, trunk_lid, wheelchair, recovery_powder, fresh_powder, framing, station, retinal_vein, fender, underbody, car, preassembly, cowl, line, enhancement_understanding, toilet, exclusive, forcel, mule, inner, bonnet, ophthalmol,	Heavy/civil construction
Logistic systems		
158	container, pallet, load, station, loading, unload, unloading, transport, lift, warehouse, pack, palletizing, stock, handle, carton, handling, empty, truck, workstation, good, delivery, palletizer, palletized, load_unload, loaded, hoist, loading_unloading, shipping, shelf, beverage,	Automated Guided Vehicles (AGVs) in manufacturing environments
180 ↑	location, navigation, identify, identification, beacon, radio, cod, map, identity, rfid_tag, global_positioning_system, radio_frequency, antenna, indoor, determine, rfid, wireless, destination, transceiver, navigate, transponder, transmitter, processor, identifier, information, system, dead_reckoning, message, include, identified,	Automated Guided Vehicles (AGVs) in manufacturing environments
Medical robotics		
47 ↑	catheter, medical, distal, patient, proximal, distal_end, sled, proximal_end, sheath, lumen, heart, elongate, bed, sterile, flexible, catheter_sheath, diagnostic_therapeutic, sterile_barrier, cardiac, elongated, handle, nose_cone, manipulation, introducer, ablation, stabilizer, sterility, steerable_catheter, therapy, ventricle,	Robot-assisted surgery and therapy
43	panoramic, imager, laparoscope, anatomical, bypass_graft, fade, zoom, artery, coronary artery, colon, secondary_containment, tract, tremor, anastomosis, swarm, insufflate_gas, heart, scissors_stapler, clamp_graspers, balloon, blood, endoscopy_laparoscopy, pavement, practitioner, prosthetist, bagger, surgery_endoscopy, file, pain, vital_sign,	Robot-assisted surgery and therapy
131 ↑	brain, tissue, tumor, neural, therapy, stereotactic, nerve, cancer, organ, mri, stimulation, biopsy, dose, clinical, radiation_therapy, lesion, compute_tomography, patient, nervous, prostate, vol, spinal_cord, sobot, emg, needle, skull, computed_tomography, multiblock, ultrasound, magnetic_resonance,	Robot-assisted surgery and therapy
110 ↑	instrument, surgical, surgeon, surgery, surgical_instrument, minimally_invasive, surgical_procedure, endoscope, surgical_site, patient, incision, tissue, medical, endoscopic, laparoscopic, medical_procedure, telesurgical, forceps, telesurgery, trocar, cannula, laparoscopic_surgery, procedure, abdominal_cavity, robotic, intuitive_surgical, abdominal_wall, console, entry, organ,	Robot-assisted surgery and therapy
126 ↑	virtual, simulation, interface, user, simulated, feedback, simulate, haptic, physical, simulator, virtual_reality, interaction, sensation, stylus, tactile, mouse, haptic_feedback, graphical, training, feel, cursor, haptic_interface, articulatable, environment, system, real, realistic, computer, experience, visual,	Other medical robots
Powered human exoskeletons		
169 ↑	joint, actuator, tendon, shoulder, actuation, muscle, wearer, rehabilitation, exoskeleton, passive, linkage, revolute_joint, degree_freedom, actuate, universal_joint, limb, ball_socket, jointed, artificial, human, elbow, mechanical, autonomous_and_remote_control_all_purpose_machine, skeletal, actuated, muscular, dofs, exercise, expandible, leg_hopping,	Powered human exoskeletons
Robots for domestic task		
72	path, movement, along, traverse, automatic_guided_vehicle, swath, follow, spatial, curvature, mow, implement, contour, refinement, optoelectronic, traversal, transformation_affine, supplementary, chess, automatic_guided_vehicles, move, mwo, centerpoint, prepositioning, desired, jolt, system, pas, matrix_mow, anticipated, ceramic_optoceramic,	Lawn-mowing
97	lens, emitter, coverage, receiver, perimeter, boundary, hydration, cliff, debris, ophthalmic_lens, lenses, eyeglass, sonar, polymerization, confinement, gateway, deterministic, monomer, piperazine, ester, airlaid_layer, pat_larsen, emission, mow, implementation, maskant, carbon_atom, spectacle_lens, acid, lawn,	Lawn-mowing
6 ↑	cleaning, clean, brush, floor, cleaner, collecting, waste, collect, carpet, scrub, collection, chassis, fan, dirty, bristle, loose_particulate, intake_port, squeegee, scrubbing, scrubber, fluid, liquid, contaminant, smear, sponge, surface, epsilon_sup, aft, dust, wet,	Vacuuming, floor cleaning
39 ↑	wheel, main_body, cleaner, chassis, dust, dock, deck, docking, docking_station, flipper, climb, caster, dust_collector, foreign_substance, payload_deck, omnidirectional_wheel, wheeled, axle, spoke, dirt, agitator, front, rotate, bristle_row, include, spacer_grid, dust_dirt, turnover, climbing, suspension,	Vacuuming, floor cleaning






Fig. C22. Word clouds of nine most central topics among 190 topics *Note:* Topics are listed in the order of centrality.



Fig. C23. Word clouds of 21 isolated topics among 190 topics.




Fig. C24. Word clouds of 190 topics.




Fig. C25. Diffusion curves for 190 topics.

References

- Acemoglu, D., Lelarge, C., Restrepo, P., 2020. Competing with robots: firm-level evidence from France. *AEA Papers and Proceedings* 110, 383–388. <https://doi.org/10.1257/pandp.20201003>.
- Acemoglu, D., Restrepo, P., 2018. Artificial intelligence, automation and work. *National Bureau of Economic Research Working Paper* 24196. <https://doi.org/10.3386/w24196>.
- Acemoglu, D., Restrepo, P., 2019. Automation and new tasks: how technology displaces and reinstates labor. *Journal of Economic Perspectives* 33 (2), 3–30. <https://doi.org/10.1257/jep.33.2.3>.
- Aghion, P., Jones, B.F., Jones, C.I., 2017. Artificial Intelligence and Economic Growth. Technical Report. NBER.
- Agrawal, A., Gans, J.S., Goldfarb, A., 2019. Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives* 33 (2), 31–50. <https://doi.org/10.1257/jep.33.2.31>.
- Alcácer, J., Gittelman, M., 2006. Patent citations as a measure of knowledge flows: the influence of examiner citations. *Review of Economics and Statistics* 88 (4), 774–779.
- Ambrosino, A., Cedrini, M., Davis, J.B., Fioria, S., Guerzoni, M., Nuccio, M., 2018. What topic modeling could reveal about the evolution of economics. *Journal of Economic Methodology* 25 (4), 367–377.
- Analytics, C., 2018. Scientific Literature and Patent Analysis of Robotics. Technical Report. Clarivate Analytics.
- Attack, J., Margo, R.A., Rhode, P.W., 2019. “Automation” of manufacturing in the late nineteenth century: the hand and machine labor study. *Journal of Economic Perspectives* 33 (2), 51–70. <https://doi.org/10.1257/jep.33.2.51>.
- Autor, D.H., Dorn, D., 2013. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review* 103 (5), 1553–1597. <https://doi.org/10.1257/aer.103.5.1553>.
- Baldwin, R., Forlids, R., 2020. Globotics and Development: When Manufacturing is Jobless and Services are Tradeable. Technical Report. NBER.
- Baumol, W., 2012. *The Cost Disease: Why Computers Get Cheaper and Health Care Doesn't*. Yale University Press, New Haven.
- Baumol, W.J., Bowen, W.G., 1966. *Performing Arts – The Economic Dilemma*. Twentieth Century Fund.
- Bessen, J., Goos, M., Salomons, A., van den Berge, W., 2020. Firm-level automation: evidence from the Netherlands. *AEA Papers and Proceedings* 110, 389–393. <https://doi.org/10.1257/pandp.20201004>.
- Bischof, J., Airolidi, E., 2012. Summarizing topical content with word frequency and exclusivity. In: Langford, J., Pineau, J. (Eds.), *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*. New York, NY: Omnipress, pp. 201–208.
- Blei, D.M., 2012. Probabilistic topic models. *Communications of the ACM* 55 (4), 77–84.
- Blei, D.M., Ng, A.Y., Jordan, M.I., 2003. Latent dirichlet allocation. *The Journal of Machine Learning Research* 3, 993–1022.
- Bouma, G., 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of the International Conference of the German Society for Computational Linguistics and Language Technology, GCSL 09*.
- Boyd-Graber, J., Hu, Y., Mimno, D., 2017. Applications of topic models. *Foundations and Trends in Information Retrieval* 11 (2–3), 143–296.
- Caselli, M., Fracasso, A., Traverso, S., 2020. Mitigation of risks of Covid-19 contagion and robotisation: Evidence from Italy. In: Press, C. (Ed.), *Covid Economics: Vetted and Real-time Papers*, pp. 174–188.
- Chae, B.K., Park, E.O., 2018. Corporate social responsibility (CSR): a survey of topics and trends using twitter data and topic modeling. *Sustainability* 10 (7), 1–20.
- Chang, H.C., 2016. The synergy of scientometric analysis and knowledge mapping with topic models: modelling the development trajectories of information security and cyber-security research. *Journal of Information & Knowledge Management* 15 (4), 1650044.
- Chen, B., Marvin, S., While, A., 2020. Containing COVID-19 in China: AI and the robotic restructuring of future cities. *Dialogues in Human Geography* forthcoming. <https://doi.org/10.1177/2043820620934267>.
- Chen, H., Zhang, G., Zhu, D., Lu, J., 2017. Topic-based technological forecasting based on patent data: a case study of Australian patents from 2000 to 2014. *Technological forecasting and social change* 119 (C), 39–52.
- Cheng, H., Jia, R., Li, D., Li, H., 2019. The rise of robots in China. *Journal of Economic Perspectives* 33 (2), 71–88. <https://doi.org/10.1257/jep.33.2.71>.
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N., 2018. Adjusting to robots: Worker-level evidence. 10.21034/iwp.13.
- Dauth, W., Findeisen, S., Suedekum, J., Woessner, N., 2019. The Adjustment of Labor Markets to Robots. Techreport. University of Würzburg. <https://sfndns.github.io/downloads/AdjustmentLaborRobots.pdf>.
- De Battisti, F., Ferrara, A., Salini, S., 2015. A decade of research in statistics: a topic model approach. *Scientometrics* 103 (2), 413–433.
- EPO, 2017. Patents and the Fourth Industrial Revolution – The Inventions behind digital transformation. Technical Report. European Patent Office.
- Eurostat, 2014. Patent Statistics: Concordance IPC V8 - NACE rev.2. Technical Report. Eurostat. https://circabc.europa.eu/sd/a/d1475596-1568-408a-9191-426629047e31/2014-10-16-Final&20IPC_NACE2_2014.pdf.
- Farrell, J., 2016. Corporate funding and ideological polarization about climate change. *Proceedings of the National Academy of Sciences* 113 (1), 92–97.
- Freeman, L.C., 1979. Centrality in social networks conceptual clarification. *Social Networks* 1 (3), 215–239.
- Goeldner, M., Herstatt, C., Tietze, F., 2015. The emergence of care robotics - a patent and publication analysis. *Technological forecasting and social change* 92, 115–131. <https://doi.org/10.1016/j.techfore.2014.09.005>.
- Graetz, G., Michaels, G., 2018. Robots at work. The review of economics and statistics 100 (5), 753–768. https://doi.org/10.1162/rest_a.00754.
- Griffith, T., Steyvers, M., 2004. Finding scientific topics. *Proceedings of the National Academy of Sciences* 101, 5228–5235.
- Grün, B., Hornik, K., 2011. Topicmodels: an R package for fitting topic models. *Journal of statistical software* 40 (13), 1–30.
- Huang, A.H., Lehavy, R., Zang, A.Y., Zheng, R., 2017. Analyst information discovery and interpretation roles: a topic modeling approach. *Management science* 64 (6), 2833–2855.
- IFR, 2018. World Robotics – Industrial Robots. Technical Report. International Federation of Robotics.
- IFR, 2018. World Robotics – Service Robots. Technical Report. International Federation of Robotics.
- IFR, 2019. IFR press conference 2019. Shanghai. <https://www.ifr.org/downloads/press/2018/IFR-World-RoboticsPresentation-18Sept2019.pdf>.
- Jaffe, A., Trajtenberg, M., 1996. Flows of knowledge from universities and federal laboratories: modeling the flow of patent citations over time and across institutional and geographic boundaries. *Proceedings of the National Academy of Sciences* 93, 12671–12677.
- Jaffe, A., Trajtenberg, M., Henderson, R., 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics* 108, 577–598.
- Kaplan, S., Vakili, K., 2015. The double-edged sword of recombination in breakthrough innovation. *Strategic Management Journal* 36, 1435–1457.
- Kim, J., Lee, J., Kim, G., Park, S., Jang, D., 2016. A hybrid method of analyzing patents for sustainable technology management in humanoid robot industry. *Sustainability* 8 (5), 1–14.
- Korzinov, V., Savin, I., 2018. General purpose technologies as an emergent property. *Technological forecasting and social change* 129, 88–104.
- Kreuchauf, F., Korzinov, V., 2017. A patent search strategy based on machine learning for the emerging field of service robotics. *Scientometrics* 111 (2), 743–772. <https://doi.org/10.1007/s11192-017-2268-3>.
- Lafond, F., Kim, D., 2019. Long-run dynamics of the U.S. patent classification system. *Journal of Evolutionary Economics* 29 (2), 631–664.
- Larsen, V.H., Thorsrud, L.A., 2019. Business Cycle Narratives. Technical Report. CESifo Working Paper Series.
- Lee, W.S., Han, E.J., Sohn, S.Y., 2015. Predicting the pattern of technology convergence using big-data technology on large-scale triadic patents. *Technological forecasting and social change* 100, 317–329.
- Lenz, D., Winker, P., 2020. Measuring the diffusion of innovations with paragraph vector topic models. *PLoS one* 15 (1), e0226685. <https://doi.org/10.1371/journal.pone.0226685>.
- Liu, L., Tang, L., Dong, W., Yao, S., Zhou, W., 2016. An overview of topic modeling and its current applications in bioinformatics. *SpringerPlus* 5 (1), 1608.
- Lüdering, J., Winker, P., 2016. Forward or backward looking? the economic discourse and the observed reality. *Journal of Economics and Statistics* 236 (4), 483–515.
- Mansfield, E., 1989. The diffusion of industrial robots in Japan and the United States. *Research policy* 18, 183–192.
- Michel, J., Bettels, B., 2001. Patent citation analysis: a closer look at the basic input data from patent search reports. *Scientometrics* 51, 185–201.
- Mimno, D., Wallach, H.M., Talley, E., Leenders, M., McCallum, A., 2011. Optimizing semantic coherence in topic models. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. PA: Association for Computational Linguistics, pp. 262–272.
- Napolitano, L., Evangelou, E., Pugliese, E., Zeppini, P., Room, G., 2018. Technology networks: the autocatalytic origins of innovation. *Royal Society open science* 5, 172445.
- Newman, D., Noh, Y., Talley, E., Karimi, S., Baldwin, T., 2010. Evaluating topic models for digital libraries. *Proceedings of the 2010 Joint Conference on Digital Libraries (JCDL)*, pp. 215–224.
- Newman, M.E.J., 2006. Finding community structure using the eigenvectors of matrices. *Physical Review E* 74, 036104.
- Pugliese, E., Cimini, G., Patelli, A., Zaccaria, A., Pietronero, L., Gabrielli, A., 2019. Unfolding the innovation system for the development of countries: co-evolution of science, technology and production. *Scientific reports* 9 (16440).
- Roberts, M.E., Stewart, B.M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S.K., Albertson, B., Rand, D.G., 2014. Structural topic models for open-ended survey responses. *American journal of political science* 58 (4), 1064–1082.
- Saracco, F., DiClemente, R., Gabrielli, A., Squartini, T., 2015. Randomizing bipartite networks: the case of the world trade web. *Scientific reports* 5, 10595.
- Savin, I., van den Bergh, J., 2021. Main topics in EIST during its first decade: a computational-linguistic analysis. *Environmental Innovation and Societal Transitions* 41, 10–17.
- Savin, I., Drews, S., van den Bergh, J., 2021. Free associations of citizens and scientists with economic and green growth: a computational linguistics analysis. *Ecological Economics* 180, 106878. <https://doi.org/10.1007/s10584-020-02842-y>.
- Savin, I., Drews, S., Maestre-Andres, S., van den Bergh, J., 2020. Public views on carbon taxation and its fairness: a computational-linguistics analysis. *Climatic change* 162, 2107–2138.
- Sievert, C., Shirley, K., 2014. LDavis: A method for visualizing and interpreting topics. *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*. Association for Computational Linguistics, Baltimore, Maryland, USA, pp. 63–70. <https://doi.org/10.3115/v1/W14-3110>. <https://www.aclweb.org/anthology/W14-3110>

- Suominen, A., Toivanen, H., Seppänen, M., 2017. Firms' knowledge profiles: mapping patent data with unsupervised learning. *Technological forecasting and social change* 115 (9), 131–142.
- Taddy, M., 2012. On estimation and selection for topic models. *Artificial Intelligence and Statistics* 1184–1193.
- Trajtenberg, M., 1990. A penny for your quotes: patent citations and the value of innovations. *RAND Journal of Economics* 21, 172–187.
- Tvinnereim, E., Fløttum, K., 2015. Explaining topic prevalence in answers to open-ended survey questions about climate change. *Nature climate change* 5, 744–747.
- Tvinnereim, E., Fløttum, K., Gjerstad, O., Johannesson, M.P., Nordø, A.D., 2017. Citizens' preferences for tackling climate change. quantitative and qualitative analyses of their freely formulated solutions. *Global Environmental Change* 46, 34–41.
- Valverde, S., Solé, R.V., Bedau, M.A., Packard, N., 2007. Topology and evolution of technology innovation networks. *Physical Review E* 76 (5), 056118.
- Venugopalan, S., Rai, V., 2015. Topic based classification and pattern identification in patents. *Technological forecasting and social change* 94, 236–250.
- WIPO, 2008. Concept of a Technology Classification for Country Comparisons. Technical Report. World Intellectual Property Organization. https://www.wipo.int/export/sites/www/ipstats/en/statistics/patents/pdf/wipo_ipc_technology.pdf
- Yang, G.-Z., Nelson, B.J., Murphy, R.R., Choset, H., Christensen, H., Collins, S.H., Dario, P., Goldberg, K., Ikuta, K., Jacobstein, N., Kragic, D., Taylor, R.H., McNutt, M., 2020. Combating COVID-19 -the role of robotics in managing public health and infectious diseases. *Robots and Society* 5 (40), eabb5589. <https://doi.org/10.1126/scirobotics.abb5589>.
- Youn, H., Strumsky, D., Bettencourt, L.M.A., Lobo, J., 2015. Invention as a combinatorial process: evidence from US patents. *Journal of The Royal Society Interface* 12 (106), 20150272.
- Zeng, Z., Chen, P.-J., Lew, A.A., 2020. From high-touch to high-tech: COVID-19 drives robotics adoption. *Tourism Geographies*. <https://doi.org/10.1080/14616688.2020.1762118>.

Ivan Savin is a senior researcher at the Institute of Environmental Science and Technology, Universitat Autònoma de Barcelona, Spain. Also Dr. Savin is a senior researcher at the Ural Federal University in Yekaterinburg, Russia. Ivan Savin completed his Ph.D. on computational econometrics at the Justus Liebig University in Giessen in 2011. In 2017 he completed his habilitation at the Karlsruhe Institute of Technology. Ivan works on topics related to economics of innovation and structural change, economics of complexity and economics of networks. More recently Ivan Savin started also working on climate policy analysis.

Ingrid Ott is a full professor at the Karlsruhe Institute of Technology (KIT) since 2010. She is leading the chair in Economic Policy. Since 2008 she is also a research associate at the Kiel Institute for the World economy (IfW). Between 2007 and 2009 she was a head of research division New Technologies and Regional Innovation Systems at the Hamburg Institute of International Economics (HWWI). And between 2002 and 2007 assistant professor at Leuphana University Lüneburg, Chair in Innovation and Growth. Her research interests are in economics of innovation and economic policy. She also has a special focus on robotics and general purpose technologies.

Chris Konop is M.Sc. from the Karlsruhe Institute of Technology (KIT). He is specializing in natural language processing.