



THE
FIFTH WAVE

BRIE-ETLA Collection of Articles

Timo Seppälä | Tomasz Mucha | Juri Mattila (eds.)

Table of contents

INTRODUCTION TO COLLECTION OF ARTICLES

**Beyond AI, Blockchain Systems, and Digital Platforms: Digitalization
Unlocks Mass Hyper-Personalization and Mass Servitization 5**

**Article 1 Estimating Firm Digitalization: A Method for Disaggregating Sector-level
Digital Intensity to Firm-level.....25**

**Article 2 AI Diffusion Monitoring among S&P500 Companies: Empirical Results
and Methodological Advancements.....49**

Article 3 What Have We Learned About Machine Learning? A Meta Analysis81

Article 4 Blockchain-Based Deployment of Product-Centric Information Systems.....99

**Article 5 Digital Protocols as Accounting and Incentivization Mechanisms in
Anti-Rival Systems: Developing a Shareable Non-Fungible Token (sNFT)139**

**Article 6 The Little Engines That Could: Game Industry Platforms and
the New Drivers of Digitalization.....161**

**Article 7 Platform-Dependent Entrepreneurs: Power Asymmetries, Risks, and
Strategies in the Platform Economy173**

**Article 8 Supranationalism, Sino-American Technology Separation, and
Semiconductors: First Observations.....211**

ARTICLE 2

AI Diffusion Monitoring among S&P500 Companies:

Empirical Results and Methodological Advancements

Tomasz Mucha^a

Timo Seppälä^{a, b}

Affiliations:

^a Aalto University, Finland

^b ETLA Economic Research, Finland

Original publication:

Mucha Tomasz and Seppala Timo, "AI Diffusion Monitoring among S&P500 Companies: Empirical Results and Methodological Advancements" (2022). DIGIT 2022 Proceedings. 13.

Abstract

With the increasing pace of digital technology innovation and commercialization, monitoring commercial diffusion of technologies becomes more important for organizations. Technology monitoring is fundamental to R&D planning, technology management, and strategic decision-making. Despite its importance, monitoring the diffusion of technologies at the commercial lifecycle stage relies on crude methods, such as “snapshot-in-time” surveys and keyword counts. These approaches are in stark contrast to novel and rapidly advancing methods for monitoring technologies at the precommercial lifecycle stages, such as fundamental scientific research and applied R&D. We address this imbalance by proposing a specialized method for monitoring the commercial diffusion of technology. The method recognizes phases in technology adoption by organizations and captures the temporal progression of the diffusion process. One of the central elements of the proposed method is the classification of text, which relies on qualitative content coding. Our approach to coding leverages the insights from innovation diffusion research and is sensitized specifically to detect phases in technology adoption by organizations. The approach is illustrated with the case of artificial intelligence (AI) diffusion among S&P 500 companies during the 2004–2019 period. Our first contribution is a new method for monitoring the commercial diffusion of technologies. It provides transparent, replicable, updatable, and granular results, which can complement survey-based technology monitoring. The second contribution is empirical evaluation of AI diffusion in the context of leading firms in North America.

Keywords

Technology diffusion, Technology adoption, Technology strategy, Artificial intelligence (AI), Machine learning (ML)

Introduction

Technology adoption is a fundamental driver of productivity and competitiveness for firms and nations (Brynjolfsson et al. 2018; Hall 2004). Hence, technology monitoring underlies the generation of strategic foresight regarding changes impacting businesses, economies, and societies (Roper et al. 2011, secs. 1 and 4.2). Therefore, monitoring technologies throughout their lifecycles is highly relevant to both research and practice. The method with the longest track record and commonly used today is survey-based research (Roper et al. 2011, pp. 100–103). Survey-based research is particularly prevalent in studies concerned with technologies entering commercialization and later stages in the technology lifecycle. Scholars and practitioners tasked with technology monitoring rely on surveys (for example, see: Balakrishnan et al. 2020; Magoulas and Swoyer 2020; Montagnier and Ek 2021; Oliveira et al. 2019; Zolas et al. 2020). Another group of technology monitoring methods, sometimes referred to as “tech mining” (Porter and Cunningham 2004), emerged from the content analysis (Roper et al. 2011, p. 106) and is currently under active development (Cunningham and Kwakkel 2016), particularly for monitoring precommercial-stage technologies. These novel approaches provide an increasing range of insights and inform R&D and technology planning related to precommercial-stage technologies.

Despite the contribution of these methods, limitations prevail in monitoring the commercial diffusion of technologies. First, since significant hurdles separate technological inventions and applied R&D from commercialization (Roper et al. 2011, p. 8), methods focused on early stages of the technology lifecycle, such as patent analysis, are not sufficient to understand the subsequent commercial diffusion of technology. Second, methods focused on later stages in the technology lifecycle also face limitations (Rogers 1983, p. 117). Thus, the development of monitoring methods suitable for commercial-stage technologies, which are longitudinal and recognize the complexity of the technology adoption process by organizations, has been missing. Therefore, we propose a method specifically designed to recognize phases in technology adoption by organizations and capture the diffusion process over time. The proposed method leverages the qualitative content analysis approach. Our approach to coding is sensitized to studying the organizational adoption of technologies. It builds on insights from innovation diffusion research concerned with the process of technology adoption within organizations (Cooper and Zmud 1990; Greenhalgh et al. 2008, sec. 5.3; Meyer and Goes 1988; Rogers 2010). The method is illustrated with the case of artificial intelligence (AI) diffusion among S&P 500 companies during the January 2004–May 2019 period. AI is a “frontier of computational advancements that references human intelligence in addressing ever more complex decision-making problems” (Berente et al. 2021, p. 1435). Top executives internationally recognize AI as having the potential to significantly impact the strategic position of their organizations and the competitive dynamics of industries (Ransboth-

am et al. 2020). Therefore, business leaders, scholars, and policy-makers are keen to monitor the commercial diffusion of AI.

This study brings several contributions. First, our method addresses the issue of technology monitoring for technologies in the latter part of their lifecycle, namely, those entering into commercialization or later stages. Second, the proposed method is versatile in terms of its applicability to a wide range of technologies. This versatility stems from its reliance on qualitative content analysis, which is not limited to any specific type of text or document, and its sensitization to broadly defined phases of technology adoption. Next, practitioners can readily adopt the proposed method into use and complement their existing technology monitoring approaches. Their projects will benefit from transparent, replicable, updatable, and granular results generated by our method. Thus, the proposed method presents a valuable addition to a survey-based approach to monitoring the commercial diffusion of technology. Finally, given that the proposed method follows a structured procedure for content coding, it may serve in the future as a foundation for an automated technology monitoring algorithm.

Theoretical background

Our approach draws on the existing research on technology monitoring and innovation diffusion. Therefore, in this section, we provide a brief overview of the relevant theory and methods from these two partially overlapping streams of literature. We separately identify the development of methods for monitoring the precommercial and commercial diffusion of technologies in both streams of literature. Technology monitoring is the process of observing and keeping up with developments in a specific technology (Roper et al. 2011, p. 72). It is widely used and provides essential inputs for both business and policy decision-makers and, thus, contributes to R&D management, technology management, and corporate and national strategies (Burgelman et al. 2004, pp. 8–9; Chen and Small 1994; Porter and Detampel 1995; Teichert and Mittermayer 2002). In this paper, we limit the scope of technology monitoring to include past developments.

Monitoring precommercial diffusion of technology

Companies cannot use precommercial-stage technologies in their daily operations but might engage with these technologies through, for example, R&D work. Nevertheless, understanding the development paths for precommercial technologies might be strategically important (Teichert and Mittermayer 2002). Since technological progress at the precommercial stage manifests itself, at least partially, in scientific

publications and patents (Porter and Cunningham 2004, p. 7), technology monitoring primarily leverages these documents (Martino 2003; Roper et al. 2011, pp. 81–82). It is often referred to as “tech mining” (Porter and Cunningham 2004, sec. 2.3; Roper et al. 2011, sec. 5.2). Porter and Cunningham (2004, p. 19) define tech mining as “the application of text mining tools to science and technology information, informed by understanding of technological innovation processes.” Particularly in the area of patent analysis, there have been many recent advances, such as analysis of innovation topics (Choi et al. 2018); identification of interindustry technologies (Fredström et al. 2021); and screening ideas in the early stages of technology development (Hong et al. 2021).

Despite these advances, monitoring the precommercial diffusion of technology is insufficient to understand the subsequent commercial diffusion. First, not all inventions “find a viable commercial application” (Grant 2016, p. 243). Next, there is a significant time lag between making an invention and its commercialization (Roper et al. 2011, sec. 1.2), which results from an innovation needing to overcome, in many cases, significant difficulties before the adopters take it into use (Rogers 2010, p. 1). Consequently, the methods for monitoring commercial diffusion of technology present a distinct area of research and practice.

Monitoring commercial diffusion of technology

Monitoring the commercial diffusion of technology develops an understanding of the extent to which the target population of adopters has taken a focal technology into use. It presents a unique set of challenges. Unlike in the case of precommercial diffusion, there are no commonly used and standardized publications to measure progression. Instead, technologies diffusing in a target market spread through various channels, such as industry conferences, press, word-of-mouth, business intelligence, and many more (Rogers 2010, pp. 18–20). Consequently, many types of actors engage in monitoring the commercial diffusion of various technologies. They include national statistical offices, not-for-profit organizations, and other service providers, such as market research firms and consulting companies.

There are two main categories of methods used in monitoring the commercial diffusion of technology: (1) survey research and (2) analysis of various types of content. We provide a brief background on the two categories and discuss their limitations. The use of surveys to collect data for research and analysis of commercial diffusion of technology has been and continues to be very prominent. A seminal study of hybrid corn diffusion in Iowa (Ryan and Gross 1943), which relied on interview-based surveys, formed the foundation of the diffusion research paradigm in the 1940s (Valente and Rogers 1995). Given the successful expansion of diffusion research in the following decades beyond the discipline of rural sociology (Rogers

2010), the previously established methodological approach continued to thrive and evolve. Recent studies investigating the diffusion of digital technologies continue to rely on surveys as a source of data (Oliveira et al. 2014, 2019). National statistical offices also use this approach to gauge the commercial diffusion of technology. A recent publication of U.S. enterprise technology adoption by the U.S. Census Bureau is a good illustration (Zolas et al. 2020). Additionally, major consulting companies and other organizations publishing insights on technology diffusion continue to rely on surveys, (for example, see: Balakrishnan et al. 2020; Magoulas and Swoyer 2020; Ransbotham et al. 2020). Despite this long lineage, survey research faces many limitations for providing insights into technology monitoring. Rogers (2010, pp. 126–130) highlights some of the criticism of survey-based methods. One of the limitations of surveys, which he points out, is providing a “snapshot-in-time” perspective rather than a “moving pictures” perspective. This low temporal granularity is a drawback, especially for rapidly advancing and diffusing technologies. Even remedying this by running surveys at multiple points in time introduces new challenges – distortion of the perception of innovation by the respondents (Rogers 1983, p. 117) and aggravation of nonresponse bias (Roper et al. 2011, p. 103). In addition, survey research in technology monitoring can suffer from long time lags, problems with definitions of technical terminology, and in the case of commercially run studies, limited transparency regarding specific methods and sampling (Montagnier and Ek 2021). Consequently, survey-based methods alone are not sufficient for monitoring the commercial diffusion of technology.

The second category of methods used in monitoring commercial diffusion of technology originates from content analysis. A study of the diffusion of multidivisional administrative structure among large industrial firms (Teece 1980) relied on qualitative content analysis. Teece analyzed, among others, annual reports, 10-K forms filed with the Securities and Exchange Commission, prospectuses, business periodical articles, recruiting literature, and publicly available texts of speeches by corporate officials. The resulting classification of organizational forms did not allow for multiple phases in innovation adoption but rather was binary in nature (Armour and Teece 1978). Similarly, a more recent analysis (Daniel Zhang et al. 2021, p. 106) disregarded phases in technology adoption by employing counts of technology-related keywords in executive presentations as an indicator of technology diffusion. Two other studies (Mikova and Sokolova 2019; Segev et al. 2015) analyzing the commercial diffusion of technology and employing content analysis also faced limitations, which resulted from the lack of control over the sample of companies included in the data analysis. Overall, we conclude that the current state of methods for monitoring the commercial diffusion of technology has been insufficient and stagnant. It is possible to address this gap by drawing on insights from innovation diffusion research on the process of technology adoption in firms (Greenhalgh et al. 2008, sec. 5.3; Rogers 2010, pp. 126–130).

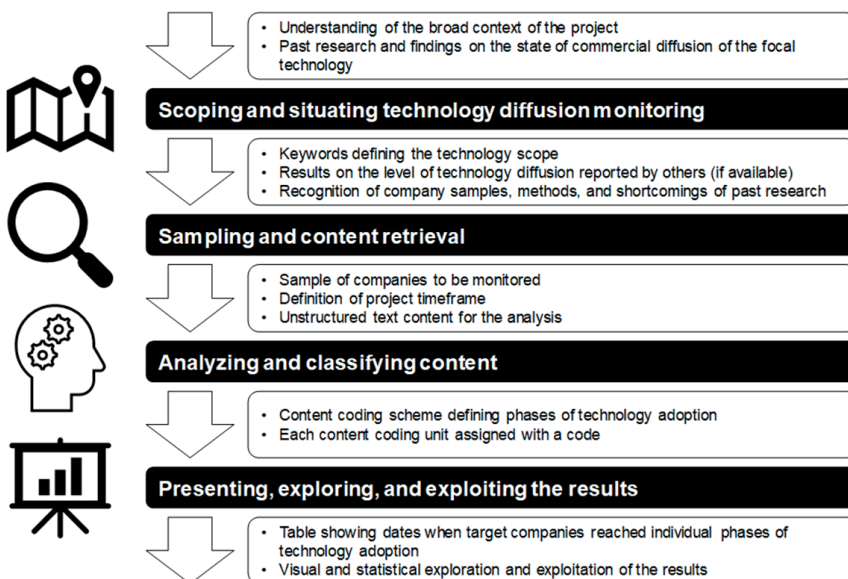
Proposed method

Our proposed approach to monitoring commercial diffusion of technology consists of four steps (Figure 1): (1) scoping and situating technology diffusion monitoring project, (2) sampling and content retrieval, (3) analyzing and classifying content, and (4) presenting, exploring, and exploiting the results. We describe these steps in greater detail in the following subsections.

Step 1: Scoping and situating technology diffusion monitoring project

The first task of researchers employing the proposed method is to define the scope of the monitoring project by identifying categories, names, or keywords representing the target technology. These keywords guide the content search and retrieval (in step 2). The terminology for describing and referring to (early) commercial-stage technologies is either established or emerging (Santos and Eisenhardt 2009). The trade-off between specificity and breadth of these keywords drives the scope of the monitoring project. For example, some keywords might represent a broader technological trajectory or frontier, such as “solar energy.” Others might encompass only a narrower set of technologies, such as “tower concentrating solar plants.” Furthermore,

Figure 1 Overview of steps in the proposed method



the level of project scoping difficulty might depend on the familiarity of the research team employing the proposed method with the target technology. If researchers are unfamiliar with the focal technology, they should first conduct a broader exploratory analysis (Roper et al. 2011, pp. 76–77).

Situating the technology diffusion monitoring project involves the identification of earlier findings on technology diffusion, which have been generated by other researchers or from commercial sources. Such reports and results on commercial diffusion of technology might be available, for example, from trade associations, market research firms, consulting firms, national statistical offices, or press. The purpose of situating technology diffusion monitoring is twofold. First, it uncovers the level of technology diffusion reported by others. These insights enable the comparison of the results from other sources against the outputs from our method (in step 4). Second, situating the project contextualizes the understanding of the diffusion process for the target technology. This understanding includes previously used definitions and scope of technology, samples of companies, methods (particularly their shortcomings), and timeframes.

Step 2: Sampling and content retrieval

Once the project scope has been defined and situated within the context of the target technology, the next step is to narrow it down and focus. This involves the selection of target companies, as well as a suitable timeframe and text content. This step concludes with the search and retrieval of unstructured text content for the analysis in the next step.

The selection of companies included in the monitored sample is vital because company size and industry are strongly associated with the rate and level of technology diffusion (Fichman 2000; Greenhalgh et al. 2008, p. 139; Oliveira and Martins 2011). Furthermore, the type of companies to be monitored will also determine the range of unstructured text sources potentially available for the analysis. Some types of content, such as websites or press articles, might be available across a wide range of companies, while larger companies might also generate content in the form of, for example, press releases, annual reports, or transcripts of executive presentations. Another aspect of content selection is its alignment between the scope of technology monitoring and the role of the technology for target companies. For example, strategically important technologies for companies in the logistics industry are likely to be discussed by these companies in press releases or annual reports, but less so in the same types of content coming from the healthcare industry, where the same technologies might still be applicable, but are not as important.

Determination of the relevant timeframe is also an integral part of this step. At the initiation of the monitoring project, it is necessary to decide how far back in time

to go. Identifying landmark events indicating technology commercialization serves that purpose well. For example, an event indicating the commercialization of wind turbines is the first installation of a utility-scale wind turbine farm by an energy company. Alternatively, patent analysis can provide insights into when a focal technology begins to enter the commercialization stage (Porter and Cunningham 2004, pp. 284–285). If, however, the monitoring project is a rerun or update of previous research, only recent information needs to be analyzed.

This step concludes with content search and retrieval. These tasks leverage technology-related keywords identified in the previous step. The content search involves the identification of documents with unstructured text content where there are references to the target technology and companies. The execution of the content search can either rely on existing commercial and open databases or custom-built approaches for content identification and retrieval. Potentially suitable content types include annual reports of listed companies, press articles, social media postings, technical reports, “gray” literature, company websites, and transcripts of executive presentations. A more in-depth discussion of the data sources, search, and content retrieval goes beyond the scope of this paper. Other authors have covered these topics in the past (for example, see Martino 2003; Mikova and Sokolova 2019; Porter and Cunningham 2004, secs. 6–8; Roper et al. 2011, sec. 5.2).

Step 3: Analyzing and classifying content

This step relies on qualitative content coding (Saldaña 2015). In the qualitative content analysis tradition, a code is “a word or short phrase that symbolically assigns a summative, salient, essence-capturing, and/or evocative attribute” to the section of text being analyzed (Saldaña 2015, p. 3). In our method, researchers generate the codes. This approach to qualitative content coding is in line with the provisional coding method, which utilizes a researcher-generated and predetermined list of codes used in the analysis (Saldaña 2015, pp. 120–123). Thus, the creation of the coding scheme (see Appendix 1) must precede the content analysis. The intention behind this scheme is to align it with the objective of the commercial technology diffusion monitoring project. Hence, the codes represent phases in the commercial adoption of technology by the target companies. Past research can provide a starting point for coding scheme development. For example, Rogers (2010) proposed a generic model describing the process of technology adoption by organizations. Appendix 2 presents a non-exhaustive list of models defining phases in innovation or technology adoption by organizations. Researchers employing the proposed method should select the initial coding scheme based on its suitability in the context of the technology monitoring project.

The analysis comprises four elements: (1) selection of the coding unit; (2) testing of the initial coding scheme on a subsample of the content; (3) potential rearrange-

ment of the scheme and another round of testing; and (4) coding of the entire sample of available content. The coding or recording unit is the “unit of text to be classified” (Weber 1990, p. 22). Since the source documents usually link to only a single company, it is sufficient to assign a single code to the whole document. However, if a single document relates to more than one company, it is necessary to narrow down the coding unit to ensure an unambiguous link between codes and individual companies. Furthermore, smaller coding units, such as paragraphs, also facilitate post-processing and post hoc analysis. For example, technology use case analysis is conducted faster when leveraging paragraph-level rather than document-level coding. After the selection of the coding unit, it is possible to test the coding scheme. Since the codes are predetermined by the researchers before analyzing the content, “[t]esting not only reveals ambiguities in the rules but also often leads to insights suggesting revisions of the classification scheme” (Weber 1990, p. 24). We suggest coding randomly selected documents representing approximately 5%–10% of the overall sample to test the coding scheme. Testing should allow researchers to evaluate whether the coding scheme granularity level is suitable. Another recommendation is to initialize the scheme with a high number of technology adoption phases. Such granularity captures finer detail from the content, if available, and thus is more informative. Researchers following this procedure must also consider reliability. There are many approaches to ensure the reliability of qualitative content coding, some of which involve quantitative measures of reliability, while others restore to consensus between raters and group discussions (Saldaña 2015, pp. 27–28). Irrespective of the selected method for establishing reliability, the researchers involved in the project should transparently report it in their study.

Step 4: Presenting, exploring, and exploiting the results

The results from the previous step need to be further processed to derive insights from technology diffusion. After coding the entire sample, the results need to be aggregated on a company and code level because it is likely that a single company will be associated with multiple documents and codes. The procedure for aggregation includes two steps: (1) sorting the documents by the company and by date from oldest to the most recent; and (2) for each company-code combination, recording the earliest date in a table. The resulting table should include company names (in rows) and phases of technology adoption included in the coding scheme (in columns). The values in the table should show dates when individual companies reached specific phases of technology adoption. Some of the cells in the resulting table are likely to be blank due to no available information. This procedure assumes that if a single company is associated with a given phase of technology adoption on a particular date, then it cannot be “degraded” to an earlier phase, even if there is a code repre-

senting a lower phase of technology adoption assigned to it on a later date. For example, researchers may code company A as reaching full-scale technology adoption in January of a given year. At the same time, based on another document from December of the same year, they may assign it a code representing testing of that technology. In that case, we assume that January is when the company has reached the full-scale commercial adoption of that technology. Researchers might still use the information about technology testing from December, for instance, in post hoc analysis (outside of the present method's scope). However, it does not impact the date of commercial adoption of the technology for company A. Thus, each date in the results table represents the earliest identified record of a given company reaching a specific phase of technology adoption.

Case Study: AI diffusion among S&P 500 companies

To illustrate the proposed method, we take the case of AI diffusion among S&P 500 companies. AI is not a single technology but rather a technological “frontier of computational advancements that references human intelligence in addressing ever more complex decision-making problems” (Berente et al. 2021, p. 1435). Technologies falling under the current umbrella of AI, most notably machine learning (ML) (Berente et al. 2021), have been recognized by executives in charge of firms around the world as having the potential to significantly impact the strategic position of their organizations and the competitive dynamics of industries (Ransbotham et al. 2020). Despite the resulting interest in AI, there is still a long way for many companies and industries to go to successfully implement the technology and have a meaningful impact on business results (Benbya et al. 2020). Hence, monitoring the progress of AI adoption by companies is a good choice for presenting the proposed method, as this technology is currently in the process of commercial diffusion among firms, particularly large firms (Benbya et al. 2020). In the remainder of this section, we present the application of the proposed method in the case of AI diffusion among S&P 500 companies.

Step 1: Scoping and situating AI diffusion

To align our search keywords with this objective, we selected “artificial intelligence” as the first target keyword. Furthermore, we recognized that the meaning of AI has been changing over recent decades (Berente et al. 2021). Therefore, we needed to limit the project scope to the latest wave of AI diffusion, which we achieved by including another broad search term representative of the current wave of AI. That second target keyword was “machine learning.” Since we were interested in diffusion across

all sectors, we did not want to favor any specific AI use case or application. Thus, we decided not to include any narrower keywords. For this method demonstration, we considered “artificial intelligence” and “machine learning” sufficient keywords to capture AI diffusion among companies.

Situating AI diffusion monitoring also involved the identification of past research and other reports on the topic. We investigated three types of sources: (1) academic research, (2) national statistical offices and other governmental or not-for-profit organizations, and (3) consulting firms and other commercially oriented organizations publishing such findings. We present an overview of key findings from each of these sources in the remainder of this subsection. Academic research concerned with or related to the diffusion of AI technologies has been expanding rapidly due to many new challenges and opportunities presented by AI (Benbya et al. 2021). Despite this interest, based on our review of the literature, scholars have largely overlooked the question of the level of AI diffusion among companies; thus, monitoring the commercial diffusion of AI has not been a focus. We have identified only a few studies that at least partially attempted to do that. In a study (Lyu and Liu 2021) investigating keywords related to AI and other technologies in job postings made by energy firms between 2010 and 2019, AI was the most common technology. It appeared in the content of 4%–8% of job postings, depending on the year. Another study (Weber and Schütte 2019) investigating AI adoption by ten globally leading retail companies analyzed content from publicly available sources generated by these companies and the press. The results indicate that eight out of ten companies leveraged AI, although there were significant differences in the level of AI infusion into the daily business operations of these companies. Finally, an annual AI Index Report (Daniel Zhang et al. 2021, p. 106) provides the absolute number of “AI” and “machine learning” mentions in corporate earnings calls. These numbers (nearly 5,000 and 1,400 mentions for AI and ML, respectively) can be compared against the historical peak of slightly above 5,000 and 2,000 mentions, respectively, and the mention counts for other technologies (which had significantly lower counts). In contrast to the limited number of studies related to AI diffusion monitoring, research giving insights into the determinants and process of AI adoption by individuals and organizations, as well as the antecedents and consequences, has been flourishing (van den Broek et al. 2020, 2021; for example, see: Grønsund and Aanestad 2020; Lebovitz et al. 2021; Lou and Wu 2021; Mayer et al. 2020; Reis et al. 2020; Strich et al. 2021; Dan Zhang et al. 2021; Zhang et al. 2020). These studies provided rich contextualization for this technology monitoring project and can inform exploration and interpretation of the results.

Understanding the diffusion of AI into commercial use by companies has been high on the agenda of many national statistical offices, government-related entities, and other not-for-profit organizations. The high priority of this topic results from the potentially high impact of AI on the economy (Ransbotham et al. 2020). The first finding that is prevalent across the results from different countries and institutions

conducting surveys is that the overall level of AI adoption is relatively low, ranging between 1% and 20% (Eurostat 2020; Montagnier and Ek 2021; Zolas et al. 2020, p. 12). Next, large organizations generally have higher adoption rates of AI than small and medium enterprises (Eurostat 2020; Montagnier and Ek 2021; Zolas et al. 2020, p. 12). There are, however, significant differences between countries. For example, the share of large enterprises with over 250 employees that analyze big data internally using machine learning is 41% for Ireland and less than 5% for countries such as Cyprus, Lithuania, Bosnia and Herzegovina (Eurostat 2020). These findings come from surveys, which suffer from limitations beyond those we discussed previously. For example, different national statistical offices rely on their own definitions of AI, thus limiting the comparability of the findings (Montagnier and Ek 2021). Some studies include multiple technologies in a basket, thus limiting the visibility of AI-only diffusion. Notably, some studies in this category employ methodologies other than surveys, namely, content analysis of company websites (Mattila et al. 2017) and patent analysis (Toole et al. 2020).

Finally, management consulting firms and other commercially oriented organizations have been the most active publishers of reports on the state of AI diffusion among companies. These reports represent the majority of the volume and variety of insights on AI diffusion out of the three types of sources we have identified. Given the sheer number of publications in this category, we concentrated on a selected few, which we considered the most representative, informative, and credible. This selectiveness means that we left out many of the reports falling into this category. We justify this decision with the significant limitations faced by publications of this type (Montagnier and Ek 2021). Frequently, the methods used were not transparent or, at least, not replicable. Since some studies sourced survey responses from proprietary contact lists (neither random nor theoretical sampling), which were undisclosed for commercial and confidentiality reasons, they were not accessible to impartial third parties. Thus, such studies were not replicable, even if they provided generic sample descriptions. These practices might lead to (un)intentional selection bias by targeting, for example, (prospective) customers with survey questionnaires. We also excluded from our analysis some reports that intentionally introduced selection bias by targeting only respondents from firms already engaged in AI activities. These reports ignored companies to which AI has not yet diffused. Finally, the commercial interests of the report writers may conflict with their readers' interests. On the positive side, these reports typically went beyond covering the state of AI diffusion and investigated topics such as related challenges faced by organizations, level of in-house expertise, numbers, type, budget, and importance of projects related to AI, roles, and count of employees involved in AI. Additionally, these reports tended to be more up-to-date than the results from academic publications or national statistical offices, given their publication volume and frequency. Overall, these reports provided us with rich insight but required careful consideration of their methods

and validity. We found that the level of AI adoption grew steadily from 2017 to 2020, with commercial AI adoption reaching 50%–60% of survey respondents or companies surveyed (Balakrishnan et al. 2020; Bughin et al. 2017; Cam et al. 2019; Chui and Malhotra 2018; Lorica and Loukides 2018; Lorica and Nathan 2019; Magoulas and Swoyer 2020; Ransbotham et al. 2017, 2018, 2019, 2020). Thus, the process of commercial diffusion of AI is still underway as we write this paper.

Step 2: Focusing on S&P 500 companies and the current wave of AI commercialization

To further narrow down the scope of the technology monitoring project, we decided to concentrate on the largest companies in a single country. We selected the largest U.S.-based companies as our target population. Based on the findings from the previous step, they were among the most advanced users of AI. Furthermore, the choice of a single country increased the homogeneity of sample companies and the content to be analyzed. These companies share an external environment and present similar internal institutional characteristics. This setting makes them sufficiently comparable to jointly analyze their commercial diffusion pattern for AI. Next, all these companies produce content in English, which allowed us to carry out the analysis in a single language only. We assumed that companies included in the S&P 500 index were representative of the target population.

Selecting the specific timeframe to be used in the analysis was the next task. Since AI has been changing the meaning over time, we wanted to exclude earlier waves of AI from the timeframe. The technology category representing the earlier wave was “expert systems” (Berente et al. 2021). We used that keyword and searched in the Scopus database for academic papers mentioning it to identify that wave. The number of articles including “expert systems” in the title, abstract, or keywords stabilized after approximately year 2000. Next, we identified the timing of several landmark events, which coincided with the start of the current commercialization wave of AI. Such events include, among others, the use of GPUs (graphics processing units) to train artificial neural networks for the first time by Andrew Ng in 2009; IBM Watson winning in Jeopardy in 2011; deep neural network-based algorithm winning the ImageNet image classification contest in 2012; and Google’s AlphaGo winning against Lee Sedol in the game of Go in 2016 (Chui et al. 2018). We decided to fix the start of the technology monitoring timeframe to January 2004, which gave five years before the first identified landmark event from the current wave of AI and four years after the number of papers related to “expert systems” stabilized. The end of the monitoring timeframe coincided with the date we retrieved the data, which was the end of May 2019.

Next, we selected the content for analysis in the technology monitoring project. Based on the findings from earlier research and reports covering AI diffusion and

use by companies, we knew that the technologies in our scope were of strategic importance. Thus, we decided to use transcripts of quarterly earnings calls and other investor presentations as content for the analysis. All sample companies were publicly listed, which meant that they all produced this type of content. Since investor events typically take the form of online conferences, detailed transcripts were available. Such events are the hallmark of voluntary disclosure (Rogers 2000) and serve two primary purposes for firms: informational and relational (Crawford Camiciottoli 2010). Tasker (1998) found that companies that provide less informative financial statements tend to make up for it with increased information content in conference calls. Additionally, the information content of the conference calls typically goes beyond the financial figures and includes forecasting and discussions on future trends, other relevant topics, and an unscripted Q&A session (Crawford Camiciottoli 2010). Thus, some investor calls include a discussion on technology development and adoption by companies. This type of content is not without limitations, such as evidence that executives engage in promotional rhetoric aimed at instilling investor confidence (Crawford Camiciottoli 2010) and may make deceptive statements (Larcker and Zakolyukina 2012). Executives might also not disclose the use of strategically important technologies. This secrecy may originate from the fiduciary responsibilities they hold toward the corporations employing them (Tiwari and Ahamed 2018) and, in some cases, personal liability. Despite these limitations, some scholars have utilized such transcripts as input data for their analysis. For example, Wang and colleagues (2020) used transcripts of earnings calls in connection with an ML-based personality trait detector to analyze executive personality impact on mergers and acquisition intensity. Teece (1980) used transcripts of speeches by corporate officials, in combination with other content, to study the diffusion of administrative innovation among large U.S. firms. Based on these findings, we concluded that transcripts of earnings calls and other investor presentations had the potential to be a suitable source of unstructured text for this method demonstration.

We retrieved 2,047 investor event transcripts of S&P 500 company executive presentations from the Thomson Reuters Eikon database. The search query was case insensitive and was “*artificial intelligence*” or “*machine learning*”. We included only events that took place between January 2004 and May 2019. Furthermore, these events were limited to quarterly earnings calls, conferences, financial analyst days, and other investor events targeting the business and investor community. The transcripts were in raw text (unstructured) format and included three metadata fields: event date, RIC (company identifier used in the database), and company name. Additionally, we collected from the same source the following data on each sample company and based on the latest available full financial year: annual revenue, primary and secondary NAICS sector codes and the respective sector names, yearly revenue per sector code (where available), and company sector based on the assignment to S&P sector indices. We used these additional data (in step 4) for the exploration and validation of the results.

Step 3: Analysis of transcripts and classification of companies into three phases of AI implementation

Before performing the analysis of content, we initialized the coding scheme based on past research. Subsequently, we tested it on a subsample of the content and revised iteratively until concluding the process with three codes: (1) mentioning AI; (2) piloting AI; and (3) commercial use of AI. Table 1 provides definitions of the codes and examples of quotes illustrating the type of statements made by company executives, which led us to assign these codes. Next, we describe in greater detail the procedure of the coding scheme development.

We initialized the coding scheme development by considering a well-known model of the information technology implementation process (Cooper and Zmud

Table 1 The final coding scheme used in the analysis of executive presentation transcripts

| Code | Definition of the code | Examples from coded texts |
|----------------------------|--|--|
| 1: Mentioning AI | <ul style="list-style-type: none"> - Reference to specific plans regarding AI or ML technology implementation - Expression of interest in or intention to implement the technology in the future - Other general reference to AI or ML | <ul style="list-style-type: none"> - “And to the extent that we can get machine learning on the volume of data that we collect, I think that’s a great opportunity for us.” - “As you would expect, head count additions primarily align with our priority areas, such as cloud and apps and machine learning.” |
| 2: Piloting AI | <ul style="list-style-type: none"> - Reference to ongoing tests, trials or experiments that involve AI or ML technologies - Any implementation of the technology that is not yet used in regular business (not part of product/offering nor regular business process) and no information regarding timing of commercial use - AI or ML related acquisition or partnership with no details on degree of commercial use of AI or ML | <ul style="list-style-type: none"> - “We’re doing a lot of work in our labs looking ahead again to the next few years in things like [...] artificial intelligence which is moving very fast [...]” - “The acquisition that we have now [...], a small company but really brings some great machine learning and vision tools [...]” - “[...] machine learning, we are [...] really prototyping that technology internally [...]” |
| 3: Commercial use of AI | <ul style="list-style-type: none"> - Reference to a current commercial use of AI or ML technologies (as part of customer offering or internal processes, which are “business as usual”) - Commercial launch or implementation utilizing AI or ML technologies in the near future (specific details provided) | <ul style="list-style-type: none"> - “We’re using software and algorithms to make decisions rather than people [...], especially as we insert machine learning into those decisions.” - “We have had great success using [...] machine-learning technologies drive those rigs to even higher levels of efficiency.” |

1990) and a classification scheme used in a practitioner-focused study investigating business adoption of AI (Ransbotham et al. 2017). The former model includes six phases: (1) initiation, (2) adoption, (3) adaptation, (4) acceptance, (5) routinization, and (6) infusion, while the latter includes five classes: (1) has not adopted AI and has no plans to do so, (2) has not adopted AI but plans to do so in the future, (3) has one or more AI pilot projects, (4) AI is incorporated in some processes and offerings, and (5) AI is extensively incorporated in processes and offerings. The two schemes present a high degree of alignment with each other. Given the focus of the latter on AI, we decided to use that scheme as our initial codes, with the addition of one code—other nonbusiness-related references to AI—to account for executives referring to AI or ML in a general sense or without giving sufficient detail on the level of commercial adoption. This initial coding scheme had a high degree of granularity and, thus, could capture a great degree of nuance in the data, if available.

Since the exclusive focus of our analysis was the identification of AI adoption phases by companies, we selected the unit of content analysis to be an individual transcript. The code assigned to each analyzed transcript corresponded to the highest degree of technology adoption identified within that transcript. This coding unit was deemed sufficient to meet the objectives of this method demonstration. Furthermore, this approach allowed us to focus the qualitative analysis only on those parts of the transcripts related to AI or ML mentions. This approach meant that for each transcript, we first identified all occurrences of relevant keywords and iteratively read paragraphs surrounding these keywords to determine sections of text that were relevant for the analysis and provided sufficient context to classify that individual document. If more than one section of text included references to AI, we coded the transcript with the highest identified level of AI adoption.

Two researchers (the first author and a research assistant knowledgeable about business use of information technology) tested the initial coding scheme. We independently coded 100 randomly selected transcripts, which represented approximately 5% of our document sample. After cross-checking the results and discussing whether the codes captured the relevant information in the transcripts, we concluded that there was a need to reduce the granularity of the coding scheme; thus, we lowered their number to four. After another round of coding, which included another set of 100 randomly selected transcripts, we cross-checked and revised the coding scheme again. The final coding scheme emerged, consisting of three codes (see Table 3). While the revision of the coding scheme aligned it better with the underlying data, this came at the expense of lower granularity, especially in the latter phases of AI adoption. This reduction in granularity points to potential limitations regarding executive transcripts as the sole content source for comprehensive technology diffusion monitoring. Nevertheless, it did not prevent us from demonstrating the proposed method and generating new insights in this case study of AI diffusion.

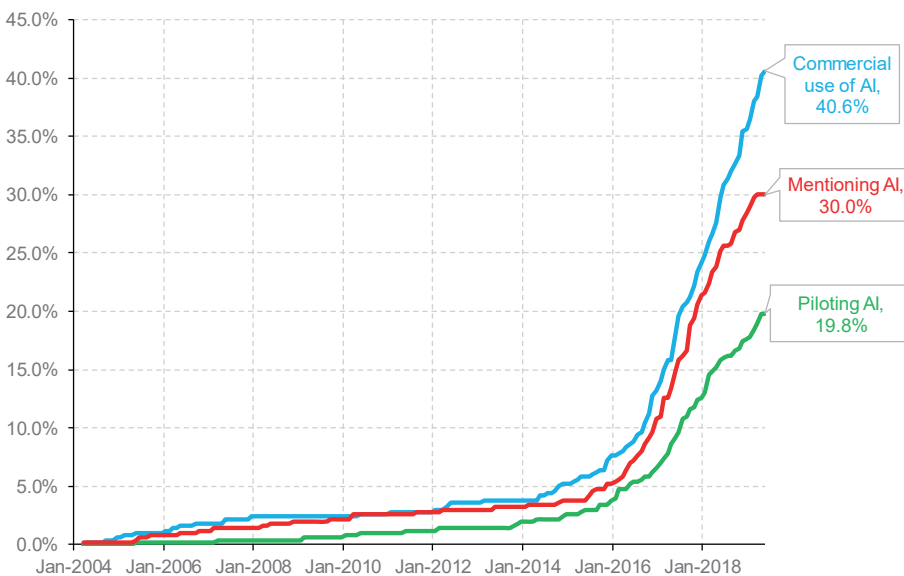
After we coded all transcripts independently, we cross-checked the results, and any differences in codes were revised and finalized through a consensus decision. According to Saldaña (2015, p. 28), this is one of the approaches used in qualitative analysis to improve consistency and address the discrepancies between coders. If our document sample was much larger, thus making parallel coding and cross-checking of the results unfeasible, or if we relied on more coders, we would restore to quantitative reliability measures.

Step 4: AI diffusion among S&P 500 companies

By aggregating the results from the previous step, which were on a document level, we arrived at the final results representing the phases of AI commercial adoption by individual S&P 500 companies throughout the monitoring timeframe. A total of 62.2% of the sample companies were assigned at least one code by the end of the study's timeframe (May 2019). As presented in Figure 2, the cumulative percentages of sample companies that reached commercial use of AI, piloted AI, and mentioned AI during investor events were 40.6%, 19.8%, and 30%, respectively.

These results are not in line with the expected sequence of technology awareness, which is followed by piloting and, later, commercial use. In other words, we expected the blue curve representing commercial use of AI to be below the two curves and

Figure 2 Cumulative percentage of companies by AI adoption phase



not above them. Based on these results, more companies reported commercial use of AI than those that either piloted AI or mentioned it in general terms during investor events. We interpret these findings as evidence of corporate executives being reluctant to build expectations by disclosing piloting of AI or referencing AI developments when their company has limited visibility on commercial implementation of AI. We conclude that the results understate the actual percentage of companies aware of AI or piloting AI. Therefore, in our subsequent discussion we primarily rely on the estimates relating to the commercial use of AI.

Validation of the results

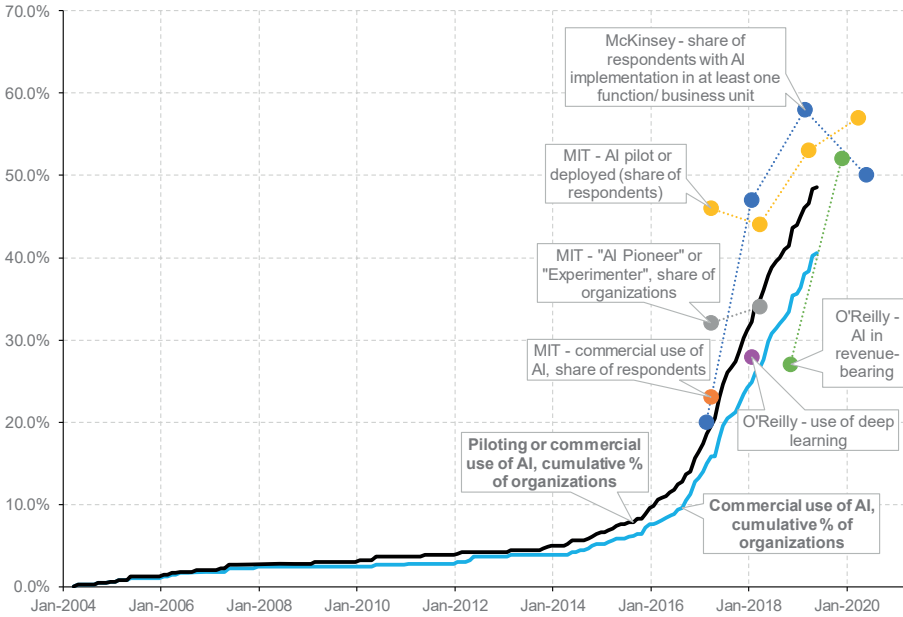
The researchers and practitioners employing the proposed method could pursue different ways of further exploring and exploiting the results presented in the previous section. We use the case study of the commercial diffusion of AI among S&P 500 companies to validate the method. We do that by comparing our results against survey-based empirical findings and two theory-based hypotheses. In this case study, our results are consistent with both empirical findings and theory.

Comparison with survey-based AI diffusion estimates

In this subsection, we compare our results on AI diffusion among S&P 500 companies with the results from several longitudinal surveys of AI use by companies, which we have identified in step 1 of the procedure. We recognize that the empirical results from these surveys are not necessarily directly comparable with our results. Thus, no formal tests can be applied here. We rely on visual inspection of Figure 3 in the results as a means of validation.

Despite limited comparability, our results on AI use by companies present an overall agreement with the trends indicated in the surveys. The alignment applies to both levels and timing. Based on this consistency, we conclude that our method provides a similar level of insight into the state of technology adoption as do commercially generated surveys. Our method, however, presents several advantages over these surveys. First, the proposed approach is transparent because it relates to a clearly defined sample of companies. The method results are also replicable due to an explicitly defined coding scheme and rules for content analysis. Another advantage of the proposed method is that the results capture a longitudinal progression of the diffusion trajectory with high granularity. What follows from transparency, replicability, and granularity is the ease of updating the results in synch with the availability of new content. Thus, the proposed method does not suffer from long time lags, which is the case with surveys. Based on the case study of AI diffusion, there appears to be

Figure 3 Comparison of the cumulative percentage of companies reaching commercial use of AI (solid blue line) and commercial use or piloting of AI (solid black line) generated using the proposed method and the results from multiple longitudinal surveys on AI use by companies



Sources of survey results: McKinsey (Balakrishnan et al. 2020; Bughin et al. 2017; Cam et al. 2019; Chui and Mathotra 2018); MIT (Ransbotham et al. 2017, 2018, 2019, 2020); O'Reilly (Lorica and Loukides 2018; Lorica and Nathan 2019; Magoulas and Swoyer 2020).

no qualitatively significant difference between the reported levels of AI use in surveys and those generated based on executive presentations geared toward investors. Thus, our method shows that it is possible to gain insight into the commercial diffusion of technology without privileged access to information using, for example, publicly available investor presentations. This result is relevant to practitioners who do not have information access similar to that of management consulting companies or other commercial organizations carrying out market analysis.

Investigation of differences between sectors

Next, we validate the results by comparing the outputs of the proposed method with the predictions generated from theory. Based on innovation diffusion theory and research results on information technology diffusion among organizations, we expect that there should be meaningful differences in the rate of AI diffusion between com-

panies from different sectors (Fichman 2000; Greenhalgh et al. 2008, p. 139; Oliveira and Martins 2011). Therefore, we can state the first null hypothesis as follows:

H1: There is no difference in the commercial diffusion rate of AI between companies from different sectors.

We can statistically test the difference between sector-level diffusion rates by investigating stochastic dominance between the diffusion curves for each sector. We examine stochastic dominance using the Kruskal–Wallis (KW) test, which is non-parametric and suitable for testing multiple groups at once (Mangiafico 2016, pp. 248–261). The test rejects the null hypothesis that there is no stochastic dominance between any pair of sectors (see Table 2). Next, we perform a post hoc analysis using the Dunn test for pairwise comparison to determine stochastic dominance individually between each pair of sectors (Mangiafico 2016, pp. 255–256). Based on this we conclude that the IT, financial, communication services, and healthcare sectors implemented AI into commercial use significantly earlier than companies in the real estate, materials, and utility sectors. These results are consistent with expectations and past empirical findings (Fichman 2000). Despite all S&P 500 companies being large corporations based in the U.S., there are meaningful differences between their commercial adoption rates of AI. In sectors where competitive pressures are highest and innovation is a driver of success, commercial AI adoption is significantly higher than in traditional sectors where fixed assets are the determinant of business success. Consequently, these results provide a validation of the proposed method against the theory.

Table 2 Results from the Kruskal–Wallis test for stochastic equality between the timing of commercial adoption of AI by different sectors

| chi-squared | df | p-value |
|-------------|----|---------|
| 87.85 | 10 | < 0.001 |

Investigation of differences between digital intensity levels

We perform another validation of the results from our methods by comparing the diffusion rates for commercial use of AI between companies exhibiting different levels of related knowledge. Related knowledge is one of the determinants that drive the adoption of information technologies by organizations (Fichman 2000; Greenhalgh et al. 2008, p. 12; Pennings and Harianto 1992). We operationalize related knowledge through the measure of the digital intensity of a firm. Digital intensity is a multifac-

eted indicator of how much firms “went digital” (Calvino et al. 2018). It measures the adoption of advanced digital technologies, employing human capital skilled with these technologies, and the extent of leveraging digital tools in relationships with customers and suppliers (Calvino et al. 2018). Based on recent empirical evidence, digital intensity is associated with AI adoption (Kinkel et al. 2021; Radhakrishnan and Chattopadhyay 2020). We use a method for approximating the digital intensity of a company based on aggregated measures of industry-level digital intensity and the level of firm engagement in different industries (Mucha and Seppälä 2021). We measure this engagement using revenue derived from activities recorded under individual business units of a company (Mucha and Seppälä 2021). Since these business units are associated with industry codes, we can map their industry-level digital intensities to the firm level. We can state the second null hypothesis as follows:

H2: There is no difference in the commercial diffusion rate of AI between companies with different levels of digital intensity.

Based on the Kruskal–Wallis (KW) test results (Table 3), we reject the null hypothesis that there is no stochastic dominance between companies from different levels of digital intensity. Post hoc analysis based on the Dunn test reveals stochastic dominance between each pair of digital intensity levels. Based on these results, we conclude that the commercial adoption of AI is strongly associated with the firm’s digital intensity level. These results are consistent with past empirical findings (Kinkel et al. 2021; Radhakrishnan and Chattopadhyay 2020), thus providing validation for the proposed method.

Table 3 Results from the Kruskal–Wallis test for stochastic equality between the timing of commercial adoption of AI by companies with different digital intensity levels

| chi-squared | df | p-value |
|-------------|----|---------|
| 54.31 | 2 | < 0.001 |

Discussion and concluding remarks

Technology monitoring is the process of observing and keeping up with developments in a specific technology (Roper et al. 2011, p. 72). It is critical to R&D management, technology management, and overall business strategy (Burgelman et al. 2004, pp. 8–9). Thus, scholars and practitioners frequently rely on technology mon-

itoring to generate new insights and knowledge. However, the predominant focus for the development of new methods for technology monitoring has been on patent analysis or otherwise precommercial stages of the technology lifecycle. These developments resulted in national statistical offices and commercial organizations relying on crude methods for monitoring the commercial diffusion of technologies, such as survey-based research developed in the 1940s and 1950s. In this paper, we propose an alternative approach to monitoring the commercial diffusion of technology.

The proposed method builds on past research within the technology monitoring and innovation diffusion literature. By utilizing qualitative content analysis, while following the procedure we propose, it is possible to generate high granularity time series representing the diffusion of technologies from early phases of commercial adoption, such as awareness of technology, to commercial use. This analysis leverages unstructured text, which can take different forms, such as the text of websites, press articles, press releases, annual reports, or transcripts of executive presentations.

We illustrate the proposed method by analyzing the commercial diffusion of AI technologies among S&P 500 companies during the January 2004–May 2019 period using 2,047 transcripts of quarterly earnings calls and other investor events. Based on qualitative content analysis of these transcripts, we assign them to one of three groups: (1) mentioning AI during investor events, (2) piloting AI, or (3) using AI in a commercial context. We find that by the end of May 2019, 40.6% of companies had reached the commercial use phase of AI, 8% reported piloting AI, and 13.6% mentioned AI in general terms only. We conclude the analysis by carrying out a validation against existing empirical findings on AI use by companies and theoretical predictions derived from the research on the diffusion of information technology among organizations. The results align well with survey results on AI diffusion published by management consulting firms and other commercially oriented organizations. Unlike these surveys, however, our method is transparent, replicable, and does not require privileged access to information, as transcripts of investor events are readily available from various databases. Another advantage of our method is that its results are available without time lags commonly associated with periodic surveys. A comparison of our results with the theoretical predictions shows consistency between the two. Our results on the differences in commercial diffusion rates for AI between companies from different sectors are consistent with expectations based on information technology diffusion research (Fichman 2000; Oliveira et al. 2019). Sectors where competitive pressure and innovativeness are high, such as IT, communication services, finance, and healthcare, adapted AI more rapidly than traditional sectors, such as utilities, real estate, and basic materials, whose fixed assets are the main determinants of competitiveness. Additionally, the results generated by our method showed that firms exhibiting a high level of digital intensity were faster commercial adopters of AI than medium or low digital intensity firms. This impact of related knowledge on the pace of AI adoption is consistent with past results from research on both in-

formation technology and AI by companies (Kinkel et al. 2021; Radhakrishnan and Chattopadhyay 2020). Overall, this illustration of the proposed method using the case of AI diffusion gives practically relevant insights and shows that the results are consistent with both past empirical findings and theoretical predictions.

This paper contributes to IS research concerned with technology monitoring and innovation diffusion as well as to practice. First, despite commercial adoption of technologies being essential to their generation of impact on economy and society (Hall 2004), this latter part of the technology life cycle has been grossly overlooked by researchers developing methods for technology monitoring. Our approach to monitoring addresses this gap by targeting technologies that enter the commercialization stage of their lifecycle or are in widespread use. Second, the proposed method is generally applicable to a wide range of technologies and contexts. This versatility results from reliance on unstructured text content as data input and broadly defined phases of technology adoption at the initiation of the analysis. Researchers employing the proposed method can fine-tune the specific content type and granularity of technology adoption phases to fit their research context. This broad applicability of the method means that it can be incorporated into and enrich a variety of studies investigating topics related to technology diffusion and adoption by organizations. These studies typically rely on surveys alone for data collection. Thus, they could increase robustness by triangulating some of the results with the method we propose. For practitioners, our method presents a transparent, replicable, and updatable alternative to commercially run surveys. Given that the proposed approach is longitudinal, ongoing technology monitoring activities carried out by strategy teams inside organizations can benefit from only incremental efforts needed to update the results with the latest analysis.

This research exhibits certain limitations and presents opportunities for further development. Since the proposed method relies on unstructured content analysis, it is of limited utility for analyzing technology diffusion among companies that generate little or no such content, such as some subpopulations of early-stage start-ups and small- and medium-sized enterprises. Next, this method might underperform surveys for studies of technologies that are of low importance to target companies. Even if the target companies generate unstructured text, mentions of such technologies might be absent there. Finally, this research relied exclusively on the case of AI diffusion among S&P 500 companies to validate the method's performance. Future validation should include a broader range of technologies, as well as types of unstructured content. Furthermore, given the continued advancements in natural language processing methods and ML, in general, the proposed method could serve as the foundation for an automated technology monitoring algorithm or tool for monitoring the commercial diffusion of technology. Such future advancement would resemble the development path of methods used in technology monitoring based on patent analysis.

References

- Armour, H. O., and Teece, D. J. 1978. "Organizational Structure and Economic Performance: A Test of the Multidivisional Hypothesis," *The Bell Journal of Economics* (9:1), p. 106. (<https://doi.org/10.2307/3003615>).
- Balakrishnan, T., Chui, M., Hall, B., and Henke, N. 2020. *Global Survey: The State of AI in 2020*, McKinsey.
- Benbya, H., Davenport, T. H., and Pachidi, S. 2020. "Special Issue Editorial: Artificial Intelligence in Organizations: Current State and Future Opportunities.," *MIS Quarterly Executive* (19:4), pp. ix-xxi.
- Benbya, H., Pachidi, S., and Jarvenpaa, S. 2021. "Special Issue Editorial: Artificial Intelligence in Organizations: Implications for Information Systems Research," *Journal of the Association for Information Systems* (22:2), Atlanta, United States: Association for Information Systems, p. 10. (<http://dx.doi.org/10.17705/1jais.00662>).
- Berente, N., Bin Gu, Recker, J., and Santhanam, R. 2021. "Managing Artificial Intelligence," *MIS Quarterly* (45:3), *MIS Quarterly*, pp. 1433-1450. (<https://doi.org/10.25300/MISQ/2021/16274>).
- van den Broek, E., Sergeeva, A., and Huysman, M. 2020. "Hiring Algorithms: An Ethnography of Fairness in Practice," in 40th International Conference on Information Systems, ICIS 2019, Association for Information Systems, pp. 1-9.
- van den Broek, E., Sergeeva, A., and Huysman, M. 2021. "When the Machine Meets the Expert: An Ethnography of Developing Ai for Hiring," *MIS Quarterly* (45:3), *MIS Quarterly*, pp. 1557-1580. (<https://doi.org/10.25300/MISQ/2021/16559>).
- Brynjolfsson, E., Rock, D., and Syverson, C. 2018. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics," in *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, pp. 23-57. (<http://www.nber.org/chapters/c14007>).
- Bughin, J., Hazan, E., Ramaswamy, S., Chui, M., Allas, T., Dahlström, P., Henke, N., and Trench, M. 2017. "Artificial Intelligence: The next Digital Frontier," McKinsey Global Institute, pp. 1-80.
- Burgelman, R. A., Christensen, C. M., and Wheelwright, S. C. 2004. *Strategic Management of Technology and Innovation*, (4th ed.), McGraw-Hill/Irwin.
- Calvino, F., Criscuolo, C., Marcolin, L., and Squicciarini, M. 2018. *A Taxonomy of Digital Intensive Sectors*. (<https://www.oecd-ilibrary.org/content/paper/f404736a-en>).
- Cam, A., Chui, M., and Hall, B. 2019. *Global AI Survey: AI Proves Its Worth, but Few Scale Impact*, McKinsey.
- Chen, I., and Small, M. 1994. "Implementing Advanced Manufacturing Technology: An Integrated Planning Model," *Omega* (22:1), pp. 91-103. ([https://doi.org/10.1016/0305-0483\(94\)90010-8](https://doi.org/10.1016/0305-0483(94)90010-8)).
- Choi, H., Oh, S., Choi, S., and Yoon, J. 2018. "Innovation Topic Analysis of Technology: The Case of Augmented Reality Patents," *IEEE Access* (6), pp. 16119-16137.

- (<https://doi.org/10.1109/ACCESS.2018.2807622>).
- Chui, M., Kamalnath, V., and McCarthy, B. 2018. "An Executive's Guide to AI," McKinsey & Company. (<https://www.mckinsey.com/~media/McKinsey/Business%20Functions/McKinsey%20Analytics/Our%20Insights/An%20executives%20guide%20to%20AI/An-executives-guide-to-AI.ashx>).
- Chui, M., and Malhotra, S. 2018. "AI Adoption Advances, but Foundational Barriers Remain," McKinsey and Company. (<https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain>).
- Cooper, R. B., and Zmud, R. W. 1990. "Information Technology Implementation Research: A Technological Diffusion Approach," *Management Science* (36:2), INFORMS: Institute for Operations Research, pp. 123–139. (<https://doi.org/10.1287/mnsc.36.2.123>).
- Crawford Camiciottoli, B. 2010. "Earnings Calls: Exploring an Emerging Financial Reporting Genre," *Discourse & Communication* (4:4), pp. 343–359. (<https://doi.org/10.1177/1750481310381681>).
- Cunningham, S. W., and Kwakkel, J. H. 2016. *Analytics and Tech Mining for Engineering Managers*, Momentum Press. (<https://books.google.fi/books?id=5yKx-DAECAAJ>).
- Eurostat. 2020. "Eurostat - Use of Machine Learning for Big Data Analysis," Eurostat. (https://ec.europa.eu/eurostat/databrowser/view/ISOC_EB_AI__custom_1269101/default/bar?lang=en, accessed September 29, 2021).
- Fichman, R. G. 2000. "The Diffusion and Assimilation of Information Technology Innovations," *Framing the Domains of IT Management: Projecting the Future through the Past* (105127), pp. 105–128.
- Fredström, A., Wincent, J., Sjödin, D., Oghazi, P., and Parida, V. 2021. "Tracking Innovation Diffusion: AI Analysis of Large-Scale Patent Data towards an Agenda for Further Research," *Technological Forecasting and Social Change* (165), p. 120524. (<https://doi.org/10.1016/j.techfore.2020.120524>).
- Grant, R. M. 2016. *Contemporary Strategy Analysis: Text and Cases Edition*, John Wiley & Sons.
- Greenhalgh, T., Robert, G., Bate, P., Macfarlane, F., and Kyriakidou, O. 2008. *Diffusion of Innovations in Health Service Organisations: A Systematic Literature Review*, John Wiley & Sons.
- Grønsund, T., and Aanestad, M. 2020. "Augmenting the Algorithm: Emerging Human-in-the-Loop Work Configurations," *The Journal of Strategic Information Systems* (29:2), *Strategic Perspectives on Digital Work and Organizational Transformation*, p. 101614. (<https://doi.org/10.1016/j.jsis.2020.101614>).
- Hall, B. H. 2004. "Innovation and Diffusion," National Bureau of Economic Research.
- Hong, S., Kim, J., Woo, H.-G., Kim, Y.-C., and Lee, C. 2021. "Screening Ideas in the Early Stages of Technology Development: A Word2vec and Convolutional Neural Network Approach," *Technovation*, p. 102407. (<https://doi.org/10.1016/j.tech>

- novation.2021.102407).
- Kinkel, S., Baumgartner, M., and Cherubini, E. 2021. “Prerequisites for the Adoption of AI Technologies in Manufacturing – Evidence from a Worldwide Sample of Manufacturing Companies,” *Technovation*, p. 102375. (<https://doi.org/10.1016/j.technovation.2021.102375>).
- Larcker, D. F., and Zakolyukina, A. A. 2012. “Detecting Deceptive Discussions in Conference Calls,” *Journal of Accounting Research* (50:2), pp. 495–540. (<https://doi.org/10.1111/j.1475-679X.2012.00450.x>).
- Lebovitz, S., Levina, N., and Lifshitz-Assaf, H. 2021. “Is Ai Ground Truth Really True? The Dangers of Training and Evaluating Ai Tools Based on Experts’ Know-What,” *MIS Quarterly* (45:3), *MIS Quarterly*, pp. 1501–1525. (<https://doi.org/10.25300/MISQ/2021/16564>).
- Lorica, B., and Loukides, M. 2018. *How Companies Are Putting AI to Work Through Deep Learning*, O’Reilly Media. Sebastopol, CA: O’Reilly Media.
- Lorica, B., and Nathan, P. 2019. *AI Adoption in the Enterprise 2019*, p. 32.
- Lou, B., and Wu, L. 2021. “Ai on Drugs: Can Artificial Intelligence Accelerate Drug Development? Evidence from a Large-Scale Examination of Bio-Pharma Firms,” *MIS Quarterly* (45:3), *MIS Quarterly*, pp. 1451–1482. (<https://doi.org/10.25300/MISQ/2021/16565>).
- Lyu, W., and Liu, J. 2021. “Artificial Intelligence and Emerging Digital Technologies in the Energy Sector,” *Applied Energy* (303), p. 117615. (<https://doi.org/10.1016/j.apenergy.2021.117615>).
- Magoulas, R., and Swoyer, S. 2020. *AI Adoption in the Enterprise 2020*, O’Reilly Media. Sebastopol, CA: O’Reilly Media.
- Mangiafico, S. S. 2016. *Summary and Analysis of Extension Program Evaluation in R*, Rutgers Cooperative Extension, New Brunswick, NJ. (<https://rcompanion.org/documents/RHandbookProgramEvaluation.pdf>).
- Martino, J. P. 2003. “A Review of Selected Recent Advances in Technological Forecasting,” *Technological Forecasting and Social Change* (70:8), *TF Highlights from ISF 2002*, pp. 719–733. ([https://doi.org/10.1016/S0040-1625\(02\)00375-X](https://doi.org/10.1016/S0040-1625(02)00375-X)).
- Mattila, J., Pajarinen, M., Rouvinen, P., and Seppälä, T. 2017. “Tekoälyn Voitto?,” *Digibarometri (2017)*, ETLA-The Research Institute of the Finnish Economy, pp. 6–30.
- Mayer, A.-S., Strich, F., and Fiedler, M. 2020. “Unintended Consequences of Introducing AI Systems for Decision Making,” *MIS Quarterly Executive* (19:4), p. 6.
- Meyer, A. D., and Goes, J. B. 1988. “Organizational Assimilation of Innovations: A Multilevel Contextual Analysis,” *Academy of Management Journal* (31:4), *Academy of Management*, pp. 897–923. (<https://doi.org/10.2307/256344>).
- Mikova, N., and Sokolova, A. 2019. “Comparing Data Sources for Identifying Technology Trends,” *Technology Analysis & Strategic Management* (31:11), *Routledge*, pp. 1353–1367. (<https://doi.org/10.1080/09537325.2019.1614157>).

- Montagnier, P., and Ek, I. 2021. "AI Measurement in ICT Usage Surveys," *OECD DIGITAL ECONOMY PAPERS*, p. 50. (<https://doi.org/10.1787/72cce754-en>).
- Mucha, T., and Seppälä, T. 2021. "Estimating Firm Digitalization: A Method for Disaggregating Sector-Level Digital Intensity to Firm-Level," *MethodsX*, p. 101233. (<https://doi.org/10.1016/j.mex.2021.101233>).
- Oliveira, T., and Martins, M. F. 2011. "Literature Review of Information Technology Adoption Models at Firm Level," *The Electronic Journal Information Systems Evaluation* (14:1), p. 13.
- Oliveira, T., Martins, R., Sarker, S., Thomas, M., and Popovic, A. 2019. "Understanding SaaS Adoption: The Moderating Impact of the Environment Context," *International Journal of Information Management* (49), pp. 1–12. (<https://doi.org/10.1016/j.ijinfomgt.2019.02.009>).
- Oliveira, T., Thomas, M., and Espadanal, M. 2014. "Assessing the Determinants of Cloud Computing Adoption: An Analysis of the Manufacturing and Services Sectors," *Information & Management* (51:5), pp. 497–510. (<https://doi.org/10.1016/j.im.2014.03.006>).
- Pennings, J. M., and Harianto, F. 1992. "The Diffusion of Technological Innovation in the Commercial Banking Industry," *Strategic Management Journal* (13:1), pp. 29–46. (<https://doi.org/10.1002/smj.4250130104>).
- Porter, A. L., and Cunningham, S. W. 2004. *Tech Mining: Exploiting New Technologies for Competitive Advantage*, (Vol. 29), John Wiley & Sons.
- Porter, A. L., and Detampel, M. J. 1995. "Technology Opportunities Analysis," *Technological Forecasting and Social Change* (49:3), pp. 237–255. ([https://doi.org/10.1016/0040-1625\(95\)00022-3](https://doi.org/10.1016/0040-1625(95)00022-3)).
- Radhakrishnan, J., and Chattopadhyay, M. 2020. "Determinants and Barriers of Artificial Intelligence Adoption – A Literature Review," in *Re-Imagining Diffusion and Adoption of Information Technology and Systems: A Continuing Conversation*, IFIP Advances in Information and Communication Technology, S. K. Sharma, Y. K. Dwivedi, B. Metri, and N. P. Rana (eds.), Cham: Springer International Publishing, pp. 89–99. (https://doi.org/10.1007/978-3-030-64849-7_9).
- Ransbotham, S., Gerbert, P., Reeves, M., Kiron, D., and Spira, M. 2018. "Artificial Intelligence in Business Gets Real," MIT Sloan Management Review and Boston Consulting Group.
- Ransbotham, S., Khodabandeh, S., Fehling, R., LaFountain, B., and Kiron, D. 2019. "Winning with AI," MIT Sloan Management Review (61180).
- Ransbotham, S., Khodabandeh, S., Kiron, D., Candelon, F., Chui, M., and LaFountain, B. 2020. "Expanding AI's Impact With Organizational Learning," MIT Sloan Management Review and Boston Consulting Group. (<https://sloanreview.mit.edu/projects/expanding-ais-impact-with-organizational-learning/>).
- Ransbotham, S., Kiron, D., Gerbert, P., and Reeves, M. 2017. "Reshaping Business with Artificial Intelligence: Closing the Gap between Ambition and Action," MIT

- Sloan Management Review (59:1), Massachusetts Institute of Technology, Cambridge, MA.
- Reis, L., Maier, C., Mattke, J., Creutzenberg, M., and Weitzel, T. 2020. "Addressing User Resistance Would Have Prevented a Healthcare AI Project Failure," *MIS Quarterly Executive* (19:4), INDIANA UNIV, OPER & DECISION TECHNOL DEPT, pp. 279–296.
- Rogers, E. M. 1983. *Diffusion of Innovations*, (3rd ed.), New York : London: Free Press ; Collier Macmillan.
- Rogers, E. M. 2010. *Diffusion of Innovations*, Simon and Schuster.
- Rogers, P. S. 2000. "CEO Presentations in Conjunction with Earnings Announcements: Extending the Construct of Organizational Genre Through Competing Values Profiling and User-Needs Analysis," *Management Communication Quarterly* (13:3), pp. 426–485. (<https://doi.org/10.1177/0893318900133003>).
- Roper, A. T., Cunningham, S., Porter, A. L., Mason, T. W., Rossini, F. A., and Banks, J. 2011. *Forecasting and Management of Technology*, (2nd ed.), Hoboken, N.J: John Wiley & Sons.
- Ryan, B., and Gross, N. C. 1943. "The Diffusion of Hybrid Seed Corn in Two Iowa Communities.," *Rural Sociology* (8:1), Rural Sociological Society, etc., p. 15.
- Saldaña, J. 2015. *The Coding Manual for Qualitative Researchers*, Sage.
- Santos, F. M., and Eisenhardt, K. M. 2009. "Constructing Markets and Shaping Boundaries: Entrepreneurial Power in Nascent Fields," *Academy of Management Journal* (52:4), Academy of Management Briarcliff Manor, NY, pp. 643–671.
- Segev, A., Jung, S., and Choi, S. 2015. "Analysis of Technology Trends Based on Diverse Data Sources," *IEEE Transactions on Services Computing* (8:6), pp. 903–915. (<https://doi.org/10.1109/TSC.2014.2338855>).
- Strich, F., Mayer, A.-S., and Fiedler, M. 2021. "What Do I Do in a World of Artificial Intelligence? Investigating the Impact of Substitutive Decision-Making AI Systems on Employees' Professional Role Identity," *Journal of the Association for Information Systems* (22:2), Atlanta, United States: Association for Information Systems, p. 9. (<http://dx.doi.org/10.17705/1jais.00663>).
- Tasker, S. C. 1998. "Bridging the Information Gap: Quarterly Conference Calls as a Medium for Voluntary Disclosure," *Review of Accounting Studies* (3:1/2), Springer Nature, pp. 137–167. (<https://doi.org/10.1023/A:1009684502135>).
- Teece, D. J. 1980. "The Diffusion of an Administrative Innovation," *Management Science* (26:5), INFORMS, pp. 464–470.
- Teichert, T., and Mittermayer, M.-A. 2002. "Text Mining for Technology Monitoring," in *IEEE International Engineering Management Conference* (Vol. 2), , August, pp. 596–601 vol.2. (<https://doi.org/10.1109/IEMC.2002.1038503>).
- Tiwari, A. K., and Ahamed, N. 2018. "Executive Tenure and Firm Performance: An Empirical Examination of Indian Corporate Landscape *," *Advances in Decision Sciences* (22), Taichung, China: Asia University, Taiwan, pp. 1–29.

- Toledo, C. 2005. "A Five-Stage Model of Computer Technology Infusion into Teacher Education Curriculum," *Contemporary Issues in Technology and Teacher Education* (5:2), Society for Information Technology & Teacher Education, pp. 177–191.
- Toole, A., Pairolero, N., Giczy, A., Forman, J., Pulliam, C., Such, M., Chaki, K., Orange, D., Homescu, A. T., and Frumkin, J. 2020. *Inventing AI: Tracing the Diffusion of Artificial Intelligence with US Patents*, October.
- Valente, T. W., and Rogers, E. M. 1995. "The Origins and Development of the Diffusion of Innovations Paradigm as an Example of Scientific Growth," *Science Communication* (16:3), pp. 242–273. (<https://doi.org/10.1177/1075547095016003002>).
- Wang, Q., Lau, R. Y. K., and Yang, K. 2020. "Does the Interplay between the Personality Traits of CEOs and CFOs Influence Corporate Mergers and Acquisitions Intensity? An Econometric Analysis with Machine Learning-Based Constructs," *Decision Support Systems* (139), p. 113424. (<https://doi.org/10.1016/j.dss.2020.113424>).
- Weber, F. D., and Schütte, R. 2019. "State-of-the-Art and Adoption of Artificial Intelligence in Retailing," *Digital Policy, Regulation and Governance* (21:3), Emerald Publishing Limited, pp. 264–279. (<https://doi.org/10.1108/DPRG-09-2018-0050>).
- Weber, R. P. 1990. *Basic Content Analysis*, Sage.
- Zhang, Daniel, Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., and Sellitto, M. 2021. "The AI Index 2021 Annual Report," ArXiv Preprint ArXiv:2103.06312.
- Zhang, Dan, Pee, L. G., and Cui, L. 2021. "Artificial Intelligence in E-Commerce Fulfillment: A Case Study of Resource Orchestration at Alibaba's Smart Warehouse," *International Journal of Information Management* (57), p. 102304. (<https://doi.org/10.1016/j.ijinfomgt.2020.102304>).
- Zhang, Z., Hummel, J. T., Nandhakumar, J., and Waardenburg, L. 2020. "Addressing the Key Challenges of Developing Machine Learning AI Systems for Knowledge-Intensive Work.," *MIS Quarterly Executive* (19:4), pp. 221–238. (<https://doi.org/10.17705/2msqe.00035>).
- Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D., Buffington, C., Goldschlag, N., Foster, L., and Dinlersoz, E. 2020. *Measuring Technology Adoption in Enterprise-Level Surveys: The Annual Business Survey*, January 3.

Appendix 1

Table 4 A generic structure of a content coding scheme

| Code | Description of the code | Examples from texts coded |
|--|--|--|
| Code 1 (Earliest considered phase of technology adoption) | <ul style="list-style-type: none"> – Provide a description of the code. – Use of negative examples (what not to include) is also useful. | <ul style="list-style-type: none"> – Provide examples (quotes) that illustrate text that should be assigned that code. – Examples might not be available in the first iteration of the coding scheme development. Therefore, in the first round of coding, coders need to rely on the code definitions alone. – This column should be populated for the subsequent coding rounds. |
| ... | ... | ... |
| Code N (Latest considered phase of technology adoption) | ... | ... |

Appendix 2

Table 5 A non-exhaustive selection of models defining phases in technology (innovation) adoption or implementation by organizations

| Source | Phases in technology (innovation) adoption or implementation by organizations |
|------------------------|--|
| (Cooper and Zmud 1990) | 1) Initiation; 2) Adoption; 3) Adaptation; 4) Acceptance; 5) Routinization; 6) Infusion |
| (Rogers 2010) | 1) Knowledge; 2) Persuasion; 3) Decision; 4) Implementation; 5) Confirmation |
| (Meyer and Goes 1988) | <p>Knowledge-awareness stage: 1) Apprehension: individuals learn of the innovation's existence; 2) Consideration: individuals consider the innovation's suitability for their organization; 3) Discussion: individuals engage in conversations concerning adoption.</p> <p>Evaluation-choice stage: 1) Acquisition proposal: it is formally proposed to purchase the equipment that embodies the innovation; 2) Medical-fiscal evaluation: medical and financial costs and benefits are weighed up; 3) Political-strategic evaluation: political and strategic costs and benefits are weighed up.</p> <p>Adoption-implementation stage: 1) Trial: the equipment is purchased but still under trial evaluation; 2) Acceptance: the equipment becomes well accepted and frequently used; 3) Expansion: the equipment is expanded or upgraded.</p> |
| (Toledo 2005) | 1) Pre-integration; 2) Transition; 3) Development; 4) Expansion; 5) Systemwide Integration |

