Tracing the Legitimacy of Artificial Intelligence –
A Media Analysis, 1980-2020

Ekaterina Korneeva
RWTH Aachen University
korneeva@time.rwth-aachen.de

Torsten Oliver Salge
RWTH Aachen University
salge@time.rwth-aachen.de

Abstract

Artificial Intelligence (AI) has received ambivalent evaluations, ranging from AI as a great opportunity and solution to crucial problems of our time to AI as a threat to humanity. For AI technologies to diffuse, they need to gain legitimacy. We trace the legitimacy of AI in society from 1980 to 2020. For our analysis, we rely on 2,543 newspaper articles from The New York Times as a reflection of societal discourse over the legitimacy of AI. Using computer-assisted content analysis, we find a sharp increase in media coverage around the mid-2010s. We find the language used in the articles to be predominantly positive and to show little changes over time. Our analysis also uncovers six highly discussed industries in the context of AI.

1. Introduction

During the last decade, Artificial Intelligence (AI), defined as intelligence performed by machines [1], has received scientific and commercial attention. The number of scientific publications, patent files, and investment has been sharply increasing since 2012 [2]. The number of patents grew by 28% annually between 2012 and 2017, with 53% of all patents on AI being published during that period [2]. Investments in AI start-ups rose from US$ 1,7 billion in 2013 to over US$ 15 billion in 2017, with a sharp increase around 2015 [3]. The developments can be traced back to an increase in computing power, a decrease in the price of computing, and the emergence of ubiquitous computing, resulting in huge amounts of stored data (big data) [3]. AI has found application in various fields. From credit scoring [4] and job applicant selection [5] automated driving [6].

Despite the achievements in this field, the development of AI and its impact on society is seen as ambivalent by scholars. Some view the immediate impacts of AI as a great opportunity for economic growth and innovation [7], while others see it as a threat. The raised concerns refer to the elimination of jobs (e.g., replacement of humans through AI-based automation) [8], the loss of privacy and self-efficacy (e.g., ubiquitous AI-technologies that trace every aspect of individuals' lives) [9], the loss of transparency of decisions [10], and possible bias and discrimination against (e.g., credit scoring and job applicant selection) [11]. In the long run, some researchers view AI as the only solution for the survival of humanity (i.e., we need AI to solve the big issues of our time) [12], while others view it as the inevitable end of humanity [e.g., as discussed by 13]. AI researchers have a very distinct and grounded opinion on AI, which they publish in research papers, conference talks, and interviews.

Likewise, users and the general public struggle with assessing the legitimacy of AI and its manifold applications. Legitimacy refers to the congruence of organizational activities and their cultural environment [14, 15]. Legitimacy is indispensable for the diffusion of novel technologies as it lays the ground for acceptance [16].

The ambivalent evaluation of AI raises questions on the legitimacy of AI in society. The developments during the last decade stress the urgency to investigate the legitimacy of AI in society. It is against this backdrop that we trace the legitimacy of AI in society from 1980 to 2020. For our analysis, we rely on mass media in reflecting public attention and evaluation of AI. We analyzed 2,543 newspaper articles from The New York Times discussing AI. We argue that public discourse reflected in media is ideally suited to shed light on the legitimacy evaluation of new, ambivalent technologies [e.g., 17].

The systematic analysis of media coverage using computer-assisted content analysis of AI allows us to understand (a) whether a public discourse is taking place, (b) in which domains (e.g., industries) it takes place, and (c) whether the discourse is rather legitimizing or de-legitimizing media framing is used. Following a qualitative approach allows us to better understand why certain developments can be observed using quantitative methods.

HICSS

Our analysis provides evidence that societal discourse follows scientific and technological developments. We observe a sharp increase in media coverage in the mid-2010s. More importantly, the language used in the articles is rather positive, indicating a legitimizing media framing, instead of a de-legitimizing media framing, dominates the media coverage. Additionally, we identified six industries that are highly discussed in the context of AI: Art and Entertainment, Automotive, Finance, Health Care, Manufacturing, and Online Services. The picture of the legitimizing media framing remained similar in the industries as in the total sample: rather positive.

2. Theoretical Background

2.1. Legitimacy

As a central concept in institutional theory, legitimacy refers to the congruence between organizational activities and their cultural environment [14, 15]. Legitimacy has been defined as "a generalized perception or assumption that the actions of an entity are desirable, proper, or appropriate within some socially constructed system of norms, values, beliefs, and definitions" (Suchman 1995, p. 574). Thereby, entities can refer to organizations, individuals, business models, industries, technologies, etc. [19].

The literature distinguishes between three kinds of legitimacy: pragmatic, cognitive, and normative [18, 19]. Pragmatic legitimacy is based on self-interested utility calculation of the organizations' immediate audience [18]. Normative legitimacy is gained by conforming to societal norms, values, and shared beliefs. Cognitive legitimacy refers to entities' comprehensiveness and take-for-grandness. While pragmatic and normative legitimacy evaluations rest on discursive evaluations, cognitive legitimacy does not [18]. The public can arrive at utility and ethical judgments through explicit public discussion. While cognitive legitimacy can be given to entities regardless of whether they have received positive, negative, or no evaluation [19]. Normative and cognitive legitimacy both refer to the wider society [16] and not only to the immediate audience, as is the case for pragmatic legitimacy [18].

Entities need to gain and maintain legitimacy to emerge, expand and survive [15]. Legitimacy has been shown to play a crucial role during the formation of new organizations [20] and new industries [19]. Emerging entities need to secure critical resources, attract investors, form demand, and convince policymakers [16, 19, 20]. Gaining legitimacy over new technologies, which are not embedded into the

social order, organizations face the "liability of newness" [21, 22], which comes with two challenges. Entities introducing a new technology face the challenge of creating new constituencies and receiving support from existing legitimate entities [18]. Ashforth and Gibbs [23] find the challenges to be especially profound for risky or uncertain technologies and the organizations' objectives are controversial. Legitimacy has been shown to positively affect the development and societal embedding of new technologies as it attracts resources from investors and support from policymakers [24]. Hence, legitimacy is especially crucial when introducing new technologies.

New technologies can gain normative legitimacy, through a public discursive evaluation [21, 22]. For societal discourse to result in legitimacy, rather than the loss of legitimacy, new technologies need to align with their environment and social rules. Only then can they pass cultural filters. Hence, collective legitimacy judgments function as a form of social control. Legitimacy is the result of a continuous societal discourse [25], which is reflected in media [26-30].

2.2. Media Coverage and Legitimacy

Media reflects and shapes public opinion and evaluation [26]. As socially embedded individuals [31], journalists integrate cultural norms and values into their work through their selection and framing of issues they report. Media coverage reflects cultural interpretation [29] and acceptability [32] of events. Hence, media reflects public approval [33]. This allows assessing collective legitimacy judgments by examining media framings.

Entman [34] states that "to frame is to select some aspects of a perceived reality and make them more salient in a communicating text, in such a way as to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation" (p. 52). Framings in media interpret certain societal developments by selectively punctuating and encoding events [35]. Media framings can approve developments while neglecting others [28]. Media framings are held together by vocabularies [17].

Relying on media coverage as a form of public discourse and legitimacy evaluation, researchers investigated the framing of different developments and events. Geels and Verhees [16] investigated the legitimacy of nuclear power from 1945 to 1986 by relying on daily newspapers. The scholars showed how the legitimacy changed over time: from initially a rather legitimizing media framing in the 1950s and 1960s to a rather de-legitimizing framing in the 1970s by referring to the anti-nuclear movement. Markard,

Wirth and Truffer [36] analyzed media framings to examine the introduction of novel technology: biogas technology. The results show how the framing of a promising technology changed from a legitimizing framing in the beginning to a de-legitimizing framing over thirteen years. Similarly, other studies relied on media coverage to study legitimacy judgments of industrial restructuring [30] and casino gambling [37]. The studies translate theoretical questions on legitimacy into empirical questions about media framings. We follow this approach and draw on media coverage studying the legitimacy of AI.

3. Research Context

Artificial Intelligence (AI) is intelligence demonstrated by machines mimicking human intelligence (Revell 2019). Intelligence is difficult to define, or as Max Tegmark puts it: "There's no agreement on what intelligence is even among intelligent intelligence researchers!" [13]. In the case of artificial intelligence, machines are considered to act intelligently when they can make the right decision under uncertainty (New Scientist 2017).

AI can manifest in computer vision, which among others, find application in autonomous vehicles, medical diagnostics, and surveillance and security. Language processing AI finds application in translation software, speech recognition, and is often used in entertainment and personal speech assistants (e.g., Apple's Siri or Amazon's Alexa). Besides the upsides, opportunities, and promises, the development of AI is also bearing some risks. The loss of jobs [e.g., 38], a threat to privacy [e.g., 39], replacing humans in general, or "breaking out" (i.e., not being in control of humans [e.g., 13], are often discussed issues in the context of AI. It is, hence, worth investigating the media framing of AI to get an understanding of public opinion on the legitimacy of this rising technology.

4. Research Method and Data

To examine the legitimacy of AI applications, we follow an established approach in communication research by drawing on newspaper articles as a reflection of public opinion [26]. In contrast to specialized magazines and niche communications (e.g., weblogs, Internet platforms), newspapers target and represent a broad audience [40]. Furthermore, renowned newspapers use objective, distancing language in their articles while abstaining from strong opinions. Unlike online sources, newspapers form a metric for established cultural associations that is

comparable over time [37]. Therefore, newspapers serve as a reliable and robust indicator of public opinion [17, 32].

4.1. Sample

For our study, we relied on newspaper articles from The New York Times (NYT). The NYT is a nationally and internationally influential daily newspaper with a high reputation and circulation [41]. The newspaper is widely perceived as an opinion leader in various fields and topics [42], including practices and implications of emerging technologies [43]. The NYT is the most cited newspaper in scientific publications, often demonstrating public discourse and interest in certain topics [44].

We retrieved NYT newspaper articles on the topic AI from 1980 to 2020 via LexisNexis. To select the relevant articles for our study from all available articles in the NYT archive, we searched for keywords in the article title and the full text. Besides "Artificial Intelligence," we also searched for closely related terms. Thereby, we relied on the systematic approach proposed by Wagener [45]. This includes a total of distinct keywords: "Artificial Intelligence," "Machine Learning," "Deep Learning," "Reinforcement Learning," "Supervised Learning," "Unsupervised Learning," and "Neural Network". If the article title or the full text entailed one of the keywords, we included it in our sample. Our search resulted in 6,811 articles.

In a second step, we inspected the frequency of keywords per article. We counted how often our keywords occurred in our selected articles. We found 4,185 articles with less than two matches with our keywords. Of those articles, we manually inspected a random sample of 100 articles. It proved that AI has not been the focus of those articles, although it was mentioned (e.g., the article was about a person winning a lottery who is also interested in AI). Hence, we dropped articles with less than two matches of our keywords. Furthermore, for the industry analysis, we relied on the indexing terms provided by LexisNexis. We dropped 83 articles due to missing indexing terms for the industry. This resulted in a sample of 2,543 NYT articles on AI with an average text length of 791 words (sd = 708, min = 24, max = 10,018).

4.2. Procedure

Our study aims to analyze the legitimacy of AI applications. For our analysis, we followed a qualitative and quantitative approach. In the first step, we employed a narrative content analysis [46]. In a second step, we relied on automated content analysis

techniques [47]. Combining quantitative and qualitative methods provides a holistic understanding of the legitimacy of AI [48]. In our qualitative analysis, we reconstructed inductively what and why particular developments of AI triggered public attention (e.g., AI beating humans in certain games). The results are presented as narratives, as descriptions of particularly distinctive events, and as an overview of fundamental developments and milestones of AI. For the quantitative analysis of public legitimacy judgment, we performed an automated content analysis of newspaper articles. To prepare our sample for the content analysis, we cleaned the newspaper articles. The cleaning process included word stemming, lemmatizing, setting to lower case, removing numbers, removing common English words (e.g., "and," "the," "or," etc.) and removing common words of our sample (e.g., "new," "york," "times").

Our aim was not only to investigate the legitimacy of AI in general but also to uncover critical industries. In the first step, we assigned industries to each article. Thereby, we relied on the industry indexing terms provided by LexisNexis for each article. The articles in our sample were assigned between 1 and 79 industry indexing terms. In total, our sample was assigned 32,626 industry indexing terms, using 500 distinct industry indexing terms. We used the North American Industry Classification System (NAICS) to classify the industry indexing terms assigned by LexisNexis. Two independent coders assigned each of the industry indexing terms to the NAICS industries. In cases where the coders assigned the indexing terms to different NAICS industries, we asked them to review the assigned indexing terms of the other coder and state whether they agree. Only if both coders agreed to assign the indexing term to a NAICS industry, we included it in our analysis. We took a threshold of a minimum of 300 articles per industry (>10% of the sample) to include the industry into our analysis. Furthermore, both coders suggested adding "Online Services" to the NAICS industries, as this was a prominent theme among the industry indexing terms assigned by LexisNexis. In total, we identified six industries that are highly discussed in the context of AI: "Arts and Entertainment," "Automotive," "Finance," "Health Care," "Manufacturing," and "Online Services".

We performed an automated sentiment analysis to systematically identify legitimizing and de-legitimizing media framings in our sample. Although unable to fully capture contextual meaning, text mining methods allow for transparency and replicability of findings. Additionally, these methods provide complete consistency in coding over the entire data set. These advantages over human assessment

make a computer-assisted content analysis of qualitative text data (i.e., media coverage) a popular approach to analyzing legitimacy judgments [49].

We relied on the established LIWC dictionaries "posemo" and "negemo" measuring positive and negative connotations for our sentiment analysis. LIWC is well-established and has been validated by psychologists, sociologists, and linguists in over a decade of work [50]. LIWC is a reliable dictionary to extract sentiment polarity from media text. Legitimizing media framings are usually associated with positively connoted language, while negatively connoted language is associated with delegitimizing media framings [36, 37, 51]. Table 1 shows an extract from the dictionary used.

Table 1: LIWC dictionary.

Sample Dictionary Words (stemmed)	
legitimizing media framing (LIWC posemo)	delegitimizing media framing (LIWC negemo)
accepta*	abuse*
brillian*	disappoint*
improve*	negclect*
opportun*	suspicio*

The sentiment of each article was determined by counting the number of positive and negative words as defined by the LIWC dictionary. This measure gives us the absolute term frequency. The absolute term frequency depends on the length of an article, which results in a bias: longer articles have a higher term frequency. To account for this, we calculated the relative term frequency. Therefore, we divided the term frequency by the total number of words of each article. In the next step, we calculated the mean relative term frequency per year to illustrate how media framing shave changed over time. Finally, we performed sentiment analyses for each industry.

5. Results

5.1. Qualitative Results

We subdivided the reconstruction of developments of AI reflected in the NYT into four phases: 1980-1989; 1990-1999, 2000-2009; 2010-2020. In this section, we provide an overview of public discourse on AI in the form of descriptive narratives [46] and draw on examples from the public discourse [52]. The findings emerge from in-depth and inductive qualitative analysis of our sample. We iterated

between data and the generated concepts until stability in interpretation was reached [53].

5.1.1. 1980-1989: Rise and Fall of AI (162 observations). In 1980 the first National Conference of the American Association for Artificial Intelligence was held at Stanford [54]. The early 1980s also marked the starting point of commercial applications based on AI research, which resulted in increased funding for academic research. However, with the collapse of the Lisp Machine market (computers optimized to process the preferred programming language for AI) by the end of the 1980s, AI lost its reputation, and a period of reduced funding and interest in AI began. This period is referred to as the second AI winter, after the first winter in the 1970s [55, 56]. Our sample provides an article from 1986 that shows high hope in computer vision for automobile manufacturing.

"After decades of research and development, a key element in the automation of American factories – devices that allow machines to 'see' – is falling into place. [...] These vision systems rely on cameras that create images, as the eye does, and on computers that interpret them, as the brain does. Proponents say manufacturers are just beginning to comprehend their vast potential." [54]

Two years later, an article titled "Setbacks for Artificial Intelligence" shows that AI's failure to fulfill the corporations' expectations of human-level intelligence and discusses the issues with Lisp.

"A major retrenchment is occurring in the artificial intelligence industry, dashing the hopes of many companies that thought they would prosper by providing the technology to make computers "think." Some of the setback stems from the failure of artificial intelligence to quickly live up to its promise [...]."[57]

"The academic researchers believed in the special programming language for artificial intelligence called Lisp, and they used computers tailored to run it. At first, sales were strong [...]. But soon that market became saturated, and artificial intelligence had trouble migrating to the mainstream of corporate America. Corporate [...] wanted to develop artificial intelligence programs without requiring their own programmers to learn Lisp." [57]

5.1.2. 1990-1999: AI can't hold up its promises (106 observations). After the second AI winter, only in the late 1990s, AI began to flourish again. An increase in computational power and the focus on specific isolated problems were named as aspects for the successful development of AI in the 90s. In 1997 the IBM chess-playing machine "Deep Blue" defeated the reigning world chess champion, Garry Kasparov [58]. However, critical voices are raised against this milestone as those claim that winning chess is not sufficient to demonstrate intelligence. In this context, the strategy game "Go" was set as the new standard to be won to demonstrate AI, while chess was put apart as primitive instead.

A repeating theme found in the 1990s is whether computers are or will be smarter than us humans. The myth that AI could mimic human intelligence led to a "crashing disappointment" in research on AI [59]. In this context, we find titles like "Smarter Than Us? Who's Us,?" "No HAL Yet: Artificial Intelligence Visions Underestimated the Mind," and "Japan Plans Computer to Mimic Human Brain".

"Deep Blue's recent trouncing of Garry Kasparov sent shock waves through the Western world. [...] Go fans proudly note, a computer has not come close to mastering what remains a uniquely human game. [...] To play a decent game of Go, a computer must be endowed with the ability to recognize subtle, complex patterns and to draw on the kind of intuitive knowledge that is the hallmark of human intelligence. 'It may be a hundred years before a computer beats humans at Go – maybe even longer,' said Dr. Piet Hut [...]. When or if a computer defeats a human Go champion, it will be a sign that artificial intelligence is truly beginning to become as good as the real thing. [58]

5.1.3. 2000-2009: The beginning of the AI spring (144 observations). The failed promises which caused the AI winter continued into the 2000s. AI had by then become a term avoided by computer scientists and software engineers "for fear of being viewed as wild-eyed dreamers" [55]. However, around the mid 2000s, there was "talk about an AI spring among researchers like Sebastian Thrun, the director of the Stanford lab" [55]. It was about the same time when DARPA introduces its Grand Challenge for autonomous vehicles and fosters research in autonomous vehicles.

"The five robots [...] demonstrated the re-emergence of artificial intelligence, a technology field that for decades has overpromised and underdelivered. [...] But the work of a small team of researchers at the Stanford Artificial Intelligence Laboratory is helping to restore credibility to the field. [...] The ability of the vehicles to complete a complex everyday task – driving – underscores how artificial intelligence may at last be moving beyond the research laboratory." [55]

Two years after a vehicle was autonomously navigated through the desert, DARPA launched the Urban Challenge for autonomous cars to obey traffic rules and drive safely in urban settings [60]. Another two years later (2009), Google, a software company – not a car manufacturer – started to build its own autonomous car [61].

5.1.4. 2010-2020: AI about to reach its potential (2131 observations). The beginning of the decade was accompanied by access to large amounts of data (i.e., big data), increasing computer power coupled with decreasing prices, and advances in machine learning technologies [62, 63]. Several events reached public's attention. In 2011, IBM's computer system, Watson, defeated the two greatest Jeopardy! champions [64]. The winning of AlphaGo against world's best Go players in 2016 and 2017 marked a milestone in the development of AI as Go was considered a complex game (e.g., compared to chess).

The year 2015 is considered a landmark for AI [65]. The use of AI at Google increased from a sporadic use in 2012 to over 2,700 projects by 2015 [65]. Not only the quantity of AI applications increased but also their performance (e.g., lower error rates in image processing) [65]. By 2016, the market for AI-related services, hardware, and software reached over 8 billion dollars [66]. Advances in big data and deep learning drove progress in video processing, text analysis, and speech recognition [67]. In 2017, 20% of companies stated to integrate "AI in some offerings or processes" [68]. In 2020, the AI lab in San Francisco "OpenAI" unveiled their new system GPT-3, a language processing system that matches human performance [69]. This achievement is seen as a breakthrough, as it comes close to human intelligence.

"For many artificial intelligence researchers, it is an unexpected step toward machines that can understand the vagaries of human language – and perhaps even tackle other human skills."[69]

5.2. Quantitative Results

To uncover the development of the public discourse of AI, we analyzed the frequency of AI-related articles in the NYT over time. An increase in the number of articles indicates an increase in public discourse. Figure 1 shows the number of articles per year. The figure shows a sharp increase in media coverage around 2015, where the number of articles triples compared to the previous year (33 articles in 2014 vs. 95 articles in 2015). In 2018 the number of

articles peaks at 552 articles. To put this into perspective, the number of articles during the first 35 years of our analysis (1980-2014) amounts to almost the same number: 553 articles. Overall, about 80% of all NYT articles published since 1980 were published from 2015-2020 (1990 articles), while only about 20% were published from 1980-2014 (553 articles).
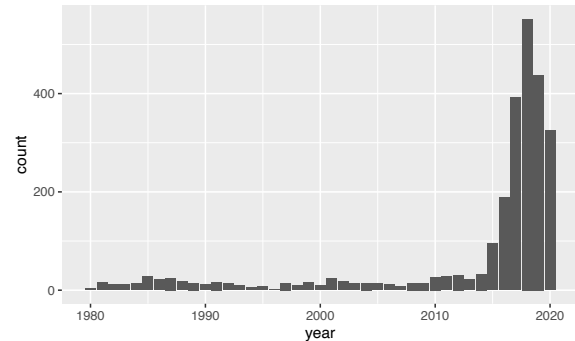

Figure 1: Number of articles per year.

To understand how the media framing has changed over time, we calculated the mean relative term frequency for positive and negative sentiment per year. Figure 2 shows a mean relative term frequency for de-legitimizing media framing (red) in 1980 of 0.032. This means that on average, 3.2% of all words used in AI-related NYT articles in 1980 were negatively connotated. While on average, around 4.5% were positively connotated (green). Overall, the relative frequency for legitimizing media framing is higher than for de-legitimizing media framing. The lines in the figure show the trends of media framing. While the red line is relatively constant, we observe a peak of the green line (legitimizing media framing) around the year 2000.
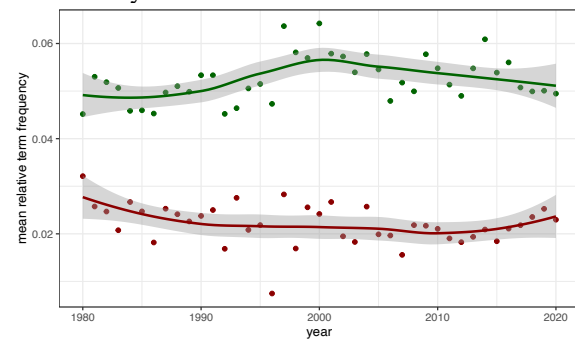

Figure 2: Relative term frequency for legitimizing (green) and de-legitimizing media framing (red).

We analyzed the industries in analogy to the analysis above. First, we counted the articles per industry per year. In a second step, we calculated the mean relative term frequency per industry per year. Figure 3 shows the results.

The industry analysis provides a similar picture as the overall analysis. We observe a sharp increase around the mid-2010s. A peak in 2018 can be observed for all but the automotive industry, which peaks a year earlier. Media coverage on online services begins in 1994, with one article and the next one being published in 1998. Although public discussion on online services starts late and slowly, online services receive the most attention according to the total number of articles (852 articles, ~1/3 of all articles). Media coverage on manufacturing and AI is considerably higher than for the other industries until 2010. For manufacturing, we can observe an increase in media coverage around 1985 and around the 2000s. Besides the number of articles, we also measured the sentiment to investigate the media framing for each industry. Figure 4 summarizes the results.

The sentiment analysis per industry pictures similar to the overall sentiment analysis. On average, the relative share of positively connoted words (mean relative term frequency) is higher than that of negatively connoted words.
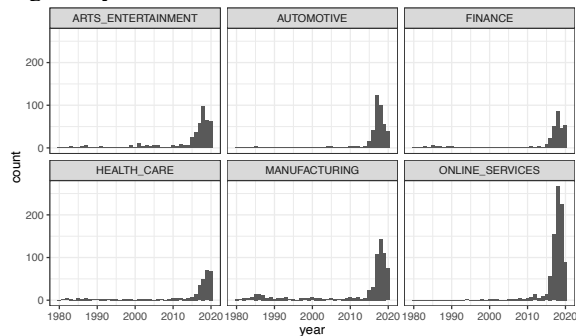

Figure 3: Number of articles per year and industry.

6. Discussion

Our results provide valuable insights into the legitimacy of AI. We observe that societal discourse, as reflected in media, follows the technological development of AI. During the AI winter, starting in the late 80s, we observe only a few articles in our sample. With technological progress, the attention increases. As AI becomes more visible, leaves the laboratory, and finds commercial applications, media attention increases. AI has received strong attention since the mid-2010s, with 80% of all articles analyzed being published since 2015. The increased media attention indicates that the legitimation process on AI is taking place [44].

Our sentiment analysis shows a higher relative term frequency for legitimizing media framing (on average, 5,2% of the words per article are positively connoted) than for de-legitimizing media framing

(on average, 2,3% of the words per article are negatively connoted). Hence, AI is discussed more often positively than negatively. This indicates that AI is perceived as an opportunity rather than a threat [70].

Looking at the sentiment over time, we observe a steady relative term frequency for positively and negatively connoted words in the articles. Especially the relative term frequency of legitimizing media framing shows little changes over time. The positive relative term frequency shows a slight increase around the year 2000. However, we do not observe overall trends (e.g., increasing or decreasing of relative term frequency) in legitimizing or de-legitimizing media framings. This finding is somehow surprising, given the increasing number of articles discussing issues on AI. We would have expected that an increase in public discourse might, over time, result in legitimizing or de-legitimizing media framing. A resulting legitimizing media framing would be visible in increasing positive relative term frequency over time or decreasing negative relative term frequency [51].

Our qualitative results provide a possible explanation for this observation. AI has repeatedly overpromised and underdelivered. Due to this history, the public might be careful putting high hopes into AI. Society might expect applications that find their way into everyday life with a visible, reliable performance by AI. Many people in the Western world have made their experiences with AI (e.g., with natural language processing in call centers or interactions with voice assistants like Alexa). Some of those experiences might have given reasons to question how intelligent AI is [71]. At the same time, society does not de-legitimize AI completely, as the potential for AI is not negatable. Our results provide evidence that we are in the middle of the legitimization process and an evaluation of legitimization has not yet taken a direction towards legitimization or de-legitimization.

Online services, as web search and social media, dominate the discussions on AI since 2015. About one-third of all articles entail online services as a theme. Our analysis identified five other highly discussed industries in the context of AI: Arts and Entertainment, Automotive, Finance, Health Care, and Manufacturing. The identification of these industries indicates that AI finds application in these, which capture society's attention. All industries follow a similar pattern: a high increase in 2015, with a peak in 2018, and a slight decrease in the last two years of our analysis. Hence, we did not find substantial differences among the industries in the media framings legitimizing or de-legitimizing AI.
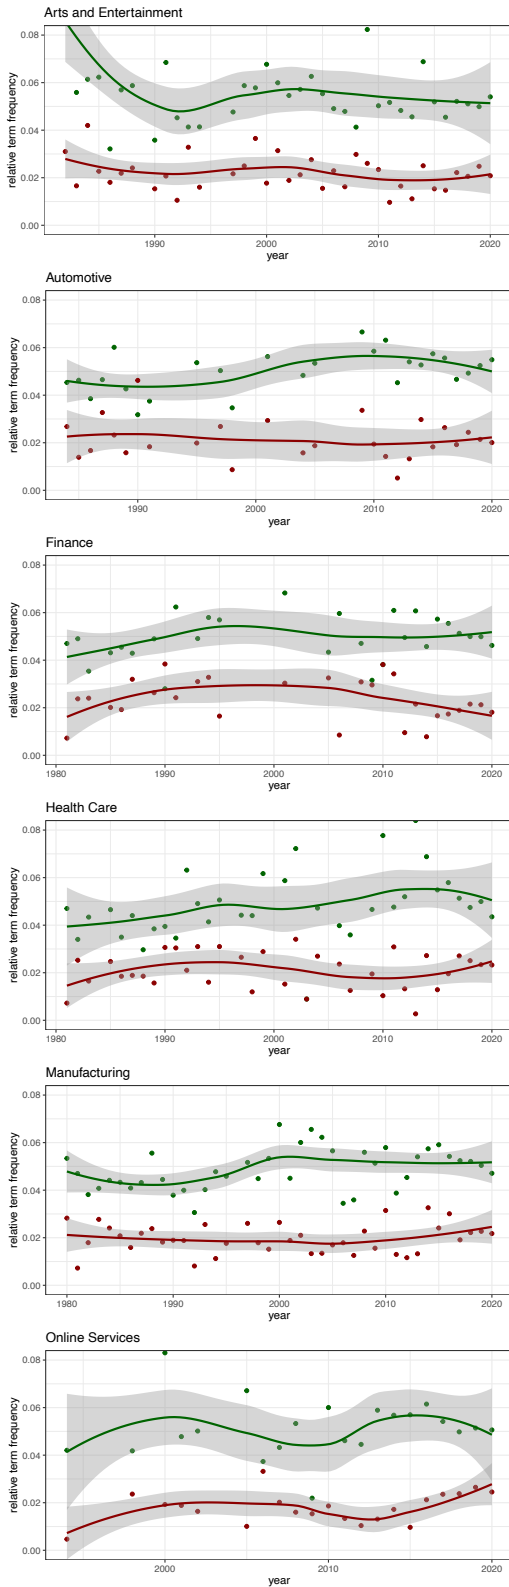
Figure 4: Relative term frequency for legitimizing (green) and de-legitimizing media framing (red) per industry.

The findings of our study provide implications for organizations and research. According to institutional theory, organizations rely on legitimacy to emerge, expand, and survive [15]. Our results indicate that there is further societal discourse needed to result in the legitimacy of AI. Previous research has investigated how organizations can foster societal discourse aiming to achieve legitimacy. Castelló, Etter and Årup Nielsen [72] have found participation in non-hierarchical open platforms and co-construction of agendas as feasible means to achieve legitimacy for sustainable development. Following a similar approach, Tegmark (2017) invites everyone to participate in the discussion on how AI will develop and how we want it to develop on the website futureoflife.org. Organizations can actively foster societal discourse for the public to evaluate the legitimacy of AI.

Our findings can be used to support empirical studies on AI acceptance and adaptation. Such studies usually draw on the technology acceptance model (TAM), measuring perceived ease of use and perceived usefulness on the individual level [71]. Our findings provide insights on the societal level and complement those findings.

7. Limitations and Future Research

Our findings are not without limitations, which provide opportunities for future research. Our analysis is based on NYT articles as a single source. This bears the risk of capturing the opinion of the readers of the media source instead of the general public. There is concern that the NYT would represent the opinion of white, educated U.S. Americans [73]. This shortcoming implicates that more media sources are needed for the robustness of our findings. Communications researchers suggest a media mix to capture public opinion (e.g., TV, social media [27].

Exploring the legitimacy of AI in other cultural regions beyond the U.S. would present another interesting extension for future research. From 2013 till 2017, 93% of all AI start-up investments went to the U.S., China, and the EU [3]. Investigating differences among these cultural regions could give insights into how factors aside from heavy investment influence the legitimacy of AI. Comparing heavily funded regions to regions with low investment in AI start-ups could give insights into the interplay of investment and legitimacy.

To develop AI applications, huge amounts of data are needed for training and later to provide services [74]. Since privacy perceptions are known to differ among cultures [75], it would be interesting to explore

the legitimacy of AI applications in regions that are known to differ in their privacy perceptions.

8. References

[1]   P. Norvig, and S. Russel, *A modern approach*: Prentice Hall Upper Saddle River, NJ, USA:, 2002.

[2]   WIPO, *Technology Trends 2019: Artificial Intelligence*, Geneva: World Intellectual Property Organization, 2019.

[3]   W. Naudé, "The race against the robots and the fallacy of the giant cheesecake: Immediate and imagined impacts of Artificial Intelligence," 2019.

[4]   R. Goh, and L. S. Lee, "Credit scoring: a review on support vector machines and metaheuristic approaches," *Advances in Operations Research,* vol. 2019, 2019.

[5]   I. Lee, and Y. J. Shin, "Machine learning for enterprises: Applications, algorithm selection, and challenges," *Business Horizons,* vol. 63, no. 2, pp. 157-170, 2020.

[6]   B. Vanholme, D. Gruyer, B. Lusetti, S. Glaser, and S. Mammar, "Highly automated driving on highways based on legal safety," *IEEE Transactions on Intelligent Transportation Systems,* vol. 14, no. 1, pp. 333-347, 2012.

[7]   K. Thomas, "Artificial Intelligence and Business Value," 2019, pp. https://towardsdatascience.com/artificial-intelligence-and-business-value-ce70083b228f.

[8]   J. Bowles, "The computerisation of European jobs," *Bruegel, Brussels*, 2014.

[9]   B. C. Stahl, and D. Wright, "Ethics and privacy in AI and big data: Implementing responsible research and innovation," *IEEE Security & Privacy,* vol. 16, no. 3, pp. 26-33, 2018.

[10]   H. Felzmann, E. F. Villaronga, C. Lutz, and A. Tamò-Larrieux, "Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns," *Big Data & Society,* vol. 6, no. 1, pp. 2053951719860542, 2019.

[11]   C. C. Perez, *Invisible women: Exposing data bias in a world designed for men*: Random House, 2019.

[12]   R. Kurzweil, *The singularity is near: When humans transcend biology*: Penguin, 2005.

[13]   M. Tegmark, *Life 3.0: Being human in the age of artificial intelligence*: Knopf, 2017.

[14]   P. J. DiMaggio, and W. W. Powell, "The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields," *American sociological review*, pp. 147-160, 1983.

[15]   J. Dowling, and J. Pfeffer, "Organizational legitimacy: Social values and organizational behavior," *Pacific sociological review,* vol. 18, no. 1, pp. 122-136, 1975.

[16]   F. W. Geels, and B. Verhees, "Cultural legitimacy and framing struggles in innovation journeys: a cultural-performative perspective and a case study of Dutch nuclear energy (1945–1986)," *Technological Forecasting and Social Change,* vol. 78, no. 6, pp. 910-930, 2011.

[17]   W. A. Gamson, and A. Modigliani, "Media discourse and public opinion on nuclear power: A constructionist approach," *American journal of sociology,* vol. 95, no. 1, pp. 1-37, 1989.

[18]   M. C. Suchman, "Managing legitimacy: Strategic and institutional approaches," *Academy of management review,* vol. 20, no. 3, pp. 571-610, 1995.

[19]   H. E. Aldrich, and C. M. Fiol, "Fools rush in? The institutional context of industry creation," *Academy of management review,* vol. 19, no. 4, pp. 645-670, 1994.

[20]   M. A. Zimmerman, and G. J. Zeitz, "Beyond survival: Achieving new venture growth by building legitimacy," *Academy of management review,* vol. 27, no. 3, pp. 414-431, 2002.

[21]   J. Freeman, G. R. Carroll, and M. T. Hannan, "The liability of newness: Age dependence in organizational death rates," *American sociological review*, pp. 692-710, 1983.

[22]   A. L. Stinchcombe, "Social structure and organizations," *Handbook of organizations*, J. G. March, ed., pp. 142-193, Chicago: Rand-McNally, 1965.

[23]   B. E. Ashforth, and B. W. Gibbs, "The double-edge of organizational legitimation," *Organization science,* vol. 1, no. 2, pp. 177-194, 1990.

[24]   M. Lounsbury, M. Ventresca, and P. M. Hirsch, "Social movements, field frames and industry emergence: a cultural–political perspective on US recycling," *Socio-economic review,* vol. 1, no. 1, pp. 71-104, 2003.

[25]   E. Vaara, and J. Tienar, "A discursive perspective on legitimation strategies in multinational corporations," *Academy of Management review,* vol. 33, no. 4, pp. 985-993, 2008.

[26]   W. A. Gamson, D. Croteau, W. Hoynes, and T. Sasson, "Media images and the social construction of reality," *Annual review of sociology,* vol. 18, no. 1, pp. 373-393, 1992.

[27]   M. Barros, "Tools of legitimacy: The case of the Petrobras corporate blog," *Organization Studies,* vol. 35, no. 8, pp. 1211-1230, 2014.

[28]   M. Joutsenvirta, and E. Vaara, "Legitimacy struggles and political corporate social responsibility in international settings: A comparative discursive analysis of a contested investment in Latin America," *Organization Studies,* vol. 36, no. 6, pp. 741-777, 2015.

[29]   R. E. Meyer, and M. A. Höllerer, "Meaning structures in a contested issue field: A topographic map of shareholder value in Austria," *Academy of Management Journal,* vol. 53, no. 6, pp. 1241-1262, 2010.

[30]   E. Vaara, J. Tienari, and J. Laurila, "Pulp and paper fiction: On the discursive legitimation of global industrial restructuring," *Organization studies,* vol. 27, no. 6, pp. 789-813, 2006.

[31]   R. D. Benford, and D. A. Snow, "Framing processes and social movements: An overview and assessment," *Annual review of sociology,* vol. 26, no. 1, pp. 611-639, 2000.

[32]   D. L. Deephouse, "Does isomorphism legitimate?," *Academy of management journal,* vol. 39, no. 4, pp. 1024-1039, 1996.

[33]   K. Lamertz, and J. A. Baum, "The legitimacy of organizational downsizing in Canada: An analysis of explanatory media accounts," *Canadian Journal of Administrative Sciences/Revue Canadienne des Sciences de l'Administration,* vol. 15, no. 1, pp. 93-107, 1998.

[34]   R. M. Entman, "Framing: Toward clarification of a fractured paradigm," *Journal of communication,* vol. 43, no. 4, pp. 51-58, 1993.

[35]   P. C. Fiss, and E. J. Zajac, "The symbolic management of strategic change: Sensegiving via framing and decoupling," *Academy of Management Journal,* vol. 49, no. 6, pp. 1173-1193, 2006.

[36]   J. Markard, S. Wirth, and B. Truffer, "Institutional dynamics and technology legitimacy–A framework and a case study on biogas technology," *Research Policy,* vol. 45, no. 1, pp. 330-344, 2016.

[37]   A. Humphreys, and K. A. Latour, "Framing the game: Assessing the impact of cultural representations on consumer perceptions of legitimacy," *Journal of Consumer Research,* vol. 40, no. 4, pp. 773-795, 2013.

[38]   J. Bessen, *AI and Jobs: The role of demand*, National Bureau of Economic Research, 2018.

[39]   H. Bae, J. Jang, D. Jung, H. Jang, H. Ha, H. Lee, and S. Yoon, "Security and privacy issues in deep learning," *arXiv preprint arXiv:1807.11655*, 2018.

[40]   G. Tuchman, "Making News: A Study in The Construction Of Reality The Free Press," *New York*, 1978.

[41]   G. Golan, "Inter-media agenda setting and global news coverage: Assessing the influence of the New York Times on three network television evening news programs," *Journalism studies,* vol. 7, no. 2, pp. 323-333, 2006.

[42]   F. Clark, and D. L. Illman, "Content analysis of New York Times coverage of space issues for the year 2000," *Science Communication,* vol. 25, no. 1, pp. 14-38, 2003.

[43]   D. Nelkin, "Science's fall from grace," *The Humanist,* vol. 55, no. 5, pp. 14, 1995.

[44]   D. Hicks, and J. Wang, "The New York Times as a resource for mode 2," *Science, technology, & human values,* vol. 38, no. 6, pp. 851-877, 2013.

[45]   A. Wagener, *Künstliche Intelligenz im Marketing-ein Crashkurs*: Haufe-Lexware, 2019.

[46]   A. Langley, "Strategies for theorizing from process data," *Academy of Management review,* vol. 24, no. 4, pp. 691-710, 1999.

[47]   A. Humphreys, *Using automated content analysis for consumer research*, working paper, December 2014, Northwestern University, Evanston, IL, 2014.

[48]   A. Tashakkori, and J. W. Creswell, "Mixed methodology across disciplines," Sage Publications Sage CA: Los Angeles, CA, 2008.

[49]   J.-P. Vergne, "Toward a new measure of organizational legitimacy: Method, validation, and illustration," *Organizational Research Methods,* vol. 14, no. 3, pp. 484-502, 2011.

[50]   J. W. Pennebaker, M. E. Francis, and R. J. Booth, "Linguistic inquiry and word count: LIWC 2001," *Mahway: Lawrence Erlbaum Associates,* vol. 71, no. 2001, pp. 2001, 2001.

[51]   P. Cichy, and T.-O. Salge, "Creating Value from Personal Data–On the Legitimacy of Business Practices in the Field of Internet-Enabled Services," 2017.

[52]   A. D. Meyer, "Adapting to environmental jolts," *Administrative science quarterly*, pp. 515-537, 1982.

[53]   K. D. Locke, *Grounded theory in management research*: Sage, 2001.

[54]   New York Times, "Factory Machines that 'see'," *March 4*, 1986.

[55]   New York Times, "Behind Artificial Intelligence, a Squadron of Bright Real People," *October 14*, 2005.

[56]   New York Times, "Brainy Robots Start Stepping Into Daily Life," *July 18*, 2006.

[57]   New York Times, "Setbacks for Artificial Intelligence," *March 4*, 1988.

[58]   New York Times, "To test a Powerful Computer Play an Ancient Game," *July 29*, 1997.

[59]   New York Times, "Deep Blue Can't Triumph in the Game of Life," *May 13*, 1997.

[60]   New York Times, "Guided by Computers and Sensors," *Ocotber 9*, 2010.

[61]   New York Times, "Google Cars Drive Themselves, in Traffic," *October 10.* , 2010.

[62]   New York Times, "As a Data Deluge Grows, Companies Rethink Storage," *March 14*, 2016.

[63]   New York Times, "Good News: A.I. Is Getting Cheaper. That's Also Bad News," *February 20*, 2018.

[64]   New York Times, "Computer Wins on 'Jeopardy!': Trivial, It's Not," *February 17*, 2011.

[65]   J. Clark, "Why 2015 Was a Breakthrough Year in Artificial Intelligence," 2015, pp. https://www.bloomberg.com/news/articles/2015-12-08/why-2015-was-a-breakthrough-year-in-artificial-intelligence.

[66]   New York Times, "IBM Is Counting on Its Bet on Watson, and Paying Big Money for It," *October 17*, 2016.

[67]   Y. Bengio, and Y. LeCun, "Scaling learning algorithms towards AI," *Large-scale kernel machines,* vol. 34, no. 5, pp. 1-41, 2007.

[68]   S. Ransbotham, D. Kiron, P. Gerbert, and M. Reeves, "Reshapinig Business with Artifical Intelligence," *MIT Sloan Mangement Review* 2017.

[69]   New York Times, "Meet GPT-3. It Has Learned to Code (and Blog and Argue)," *November 24*, 2020.

[70]   A. R. Schuck, and C. H. De Vreese, "Between risk and opportunity: News framing and its effects on public support for EU enlargement," *European journal of communication,* vol. 21, no. 1, pp. 5-32, 2006.

[71]   K. Sohn, and O. Kwon, "Technology acceptance theories and factors influencing artificial intelligence-based intelligent products," *Telematics and Informatics,* vol. 47, pp. 101324, 2020.

[72]   I. Castelló, M. Etter, and F. Årup Nielsen, "Strategies of legitimacy through social media: The networked strategy," *Journal of management studies,* vol. 53, no. 3, pp. 402-432, 2016.

[73]   D. H. Weaver, R. A. Beam, B. J. Brownlee, P. S. Voakes, and G. C. Wilhoit, *The American journalist in the 21st century: US news people at the dawn of a new millennium*: Routledge, 2009.

[74]   D. E. O'Leary, "Embedding AI and crowdsourcing in the big data lake," *IEEE Intelligent Systems,* vol. 29, no. 5, pp. 70-73, 2014.

[75]   F. Bélanger, and R. E. Crossler, "Privacy in the digital age: a review of information privacy research in information systems," *MIS quarterly*, pp. 1017-1041, 2011.