

“Mixture of amazement at the potential of this technology and concern about possible pitfalls”: Public sentiment towards AI in 15 countries

Patrick Gage Kelley,* Yongwei Yang,* Courtney Heldreth,*
Christopher Moessner,† Aaron Sedley,* Allison Woodruff*

**Google*

patrickgage@acm.org, yongwei@google.com, cheldreth@google.com,
asedley@google.com, woodruff@acm.org

†*Ipsos*

christopher.moessner@ipsos.com

Abstract

Public opinion plays an important role in the development of technology, influencing product adoption, commercial development, research funding, career choices, and regulation. In this paper we present results of an in-depth survey of public opinion of artificial intelligence (AI) conducted with over 17,000 respondents spanning fifteen countries and six continents. Our analysis of open-ended responses regarding sentiment towards AI revealed four key themes (exciting, useful, worrying, and futuristic) which appear to varying degrees in different countries. These sentiments, and their relative prevalence, may inform how the public influences the development of AI.

1 Introduction

Increased understanding of the societal impact of artificial intelligence (AI) has spurred strong interest its in responsible development [20, 30, 34, 59]. Researchers, advocates, companies, and others have proposed processes, principles, design toolkits, and other resources to support thoughtful development of AI that carefully considers both benefits and risks [1, 26, 31, 38, 27, 52, 15, 35].

Public opinion is an important force in responsible development, exerting pressure on funding agencies, regulators, companies, educators, and others to address both general attitudes and specific issues [10, 12, 51, 67], such as the impact of automation on the future of work [7, 54, 57], the interaction of AI with human rights issues such as privacy and discrimination [1, 3, 9, 16], the ethics of autonomous weapons [58, 64], and the development and availability of dual-use technologies such as synthetic media that may be used for either benevolent or nefarious purposes [50]. While public opinion may not fully align with expert assessment on these issues, it is nonetheless useful to elucidate the forces in effect.

Copyright 2021 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Bulletin of the IEEE Computer Society Technical Committee on Data Engineering

While there have been some explorations of public perception of AI, for example, survey research [2, 6, 11, 37, 45, 49, 64, 68], sentiment analysis [18, 25, 29, 51], and narrative analysis [12, 13], much of this work has been done in Western, English-speaking contexts. Even in these better studied contexts, much remains to be learned, as both the technology and the public discussion are evolving rapidly. In this paper, we present a survey of public perception of AI conducted with over 17,000 respondents spanning fifteen countries and six continents (encompassing in total: Germany, Australia, Finland, Singapore, Belgium, Canada, the United States (US), South Korea, Spain, France, Poland, Brazil, China, India, and Nigeria). Using an inductive approach to analyze open-ended responses, we identified four key sentiment groups (exciting, useful, worrying, and futuristic) whose prevalence distinguishes responses to AI in different countries. We previously shared results from eight of these countries [40] and here we extend our analysis to fifteen countries and more fully discuss the sentiments. We then discuss implications of these findings for the development of AI systems.

2 Background

Artificial intelligence (AI) is a broad term with no consensus definition [23, 25, 59], and the scope of our inquiry is intended to be similarly broad. We note that interpretation of the term is further confounded by the “AI effect” (the phenomenon that once AI successfully solves a problem and the solution becomes commonplace, it is no longer considered to be AI) [42], as well as lack of awareness of algorithmic processing in common systems [24, 53, 63]. To aid comparison with survey responses, following [59], we share with the reader the following definition provided by Nils J. Nilsson: “Artificial intelligence is an activity devoted to making machines intelligent, and intelligence is the quality that enables an entity to function appropriately and with foresight in its environment.” [48]

2.1 Empirical Studies

Much of the research on public perception of AI has been survey-based, often conducted in Western, English-speaking countries such as the US and the UK [6, 11, 23, 68, 49] although this has been broadening recently. AI is often viewed as likely to have a significant impact on the future, with a frequent expectation that its effects will be positive. In a 2019 Edelman survey in the US, 9 out of 10 respondents assumed that AI will be life-changing and transformational [23]. A Gallup survey conducted in the US in 2018 found that 76% believed that AI will have a positive impact on their lives [49]; 61% of respondents had a positive view of AI and robots in a large-scale 2017 survey across Europe on the impact of digitization and automation on daily life [60]; and a 2017 consumer research survey conducted across North America, Europe, and Asia revealed a predominant expectation that society will become better (61%) rather than worse (22%) due to increased automation and AI [2]. A recent Pew Research survey conducted across the Americas, Europe, and Asia showed a somewhat narrower margin (possibly due to shifting public opinion, or alternatively, methodological differences), with a median of 53% saying that AI has been mostly good for society (53%) versus mostly bad (33%) [28]. Considering expected impact in the next 20 years, the 2019 World Risk Poll indicated AI would mostly help (41%) versus mostly harm (30%) people in one’s own country, with more favorable impressions in Asia and less favorable impressions in Western countries [41, 47].

At the same time, AI is neither interpreted as exclusively beneficial nor exclusively disadvantageous, and public response often indicates contradictory emotions. Looking at broad reactions, Blumberg reported that US respondents were equally split between feeling optimistic and informed and feeling fearful and uninformed about AI [6], while [2] also revealed both excitement and concern. Relatedly, a 2019 Mozilla survey open to respondents on the Internet gathered continent-level demographic data and revealed varying and mixed emotions at the continent-level [45]. Specific concerns have been expressed regarding social issues, such as AI benefiting the wealthy and harming the poor, fear that AI-enabled deepfakes will erode trust in information, and AI increasing

social isolation and reducing human capability [23]. In line with these concerns, Zhang and Dafoe found that 82% of Americans want AI and robots to be carefully managed [68], with 88% of Europeans expressing similar sentiment [60]. Moreover, 60% of the general population in the Edelman survey expressed the need for more regulation regarding AI development and deployment [23].

Qualitative work has also explored public perception of algorithmic systems, for example, finding that perception of algorithmic systems can vary substantially by individual factors as well as platform [19], and that end users often have fundamental questions or misconceptions about technical details of their operation [8, 24, 53, 62, 63].

2.2 Narratives and Media Sentiment Analysis

AI is not only heavily discussed in academia, but is also a popular topic in public media and entertainment [23]. In fact, Cave et al. provide a history of narratives about intelligent machines dating back to ancient Greece [13]. In modern times, 58% of the respondents in a recent Blumberg survey indicated that they get information about AI from movies, TV, and social media [6]. In a 2016 CBS news survey, only 19% indicated not having seen any of several AI movies such as “The Terminator” or “I, Robot” [14]. Cave et al. argue that prevalent AI narratives in the English-speaking West share “a tendency towards utopian or dystopian extremes,” cautioning that inaccurate narratives could affect technological advancement and regulation [12], with similar points raised in [34, 59, 66]. Cave et al. surveyed UK respondents regarding their responses to eight dominant narratives about AI, reporting that the strong majority elicited more concern than excitement [11].

At the same time, while some researchers have argued that narratives and fiction may be disproportionately frightening, studies have suggested that news reports may be more balanced or appropriately critical. Sentiment analysis of newspaper articles from the New York Times and associated content found that, in general, AI has had consistently more optimistic than pessimistic coverage over time [25], and did not support the hypothesis that news media coverage of AI is negative [29]. Content analysis of coverage of AI in five major American newspapers revealed benefits were discussed more frequently than risks, although risks were discussed with greater specificity [18]. Ouchchy et al. analyzed discussion of AI ethics in English language media sources and concluded that “The issues most frequently covered, along with the mostly balanced/neutral tones, suggest that the media has a fairly realistic and practical focus in its coverage of the ethics of AI.” [51]

2.3 National Considerations

A number of countries have established national strategies to promote the use and development of AI, which vary by country and may influence public perception [22].¹ The importance of studying local context is also illustrated by analysis of country-specific opportunities and challenges for AI, e.g. [39]. Further, researchers have called for better integration of developing country considerations in the discussion and development of AI [56].

2.4 Our Approach

Our work sits within a growing body of research on people’s perceptions of AI, across disciplines including HCI, critical studies, law, marketing, policy, psychology, and more. This topic is highly complex, multi-dimensional, and far from fully understood. Methodologically, this means that techniques such as *triangulation* (studying the same phenomenon from multiple vantage points, in order to cross-check and more fully capture richness and complexity, e.g. using both qualitative and quantitative methods to see if the findings are consistent) [55] and *replication* (the reproduction and extension of prior work) [65] are particularly useful for this topic. Accordingly, we seek to broaden and enrich the understanding of sentiment towards AI by looking for emergent themes in a large number of open-ended responses from a wide range of countries.

¹See also <https://futureoflife.org/national-international-ai-strategies/>

3 Methodology

In order to better understand public perception of AI, we partnered with Ipsos, a global market research firm, to conduct a survey of 17,014 respondents in fifteen countries in July and December 2019.² Methodologically, this work falls in the genre of public opinion polling, as described below.

3.1 Instrument Development and Translation

To develop concepts and questions, we consulted experts at our institutions, reviewed published work, drew on our own previous unpublished research, and conducted an initial pilot survey in June 2018 with 1300 respondents drawn from a panel of the general online population in the US. Many questions in the final instrument were written uniquely for this survey while others were modified from or replicate other questions in the literature or the canon of public opinion surveys. In order to more accurately reflect real-world settings, we did not define AI, and left interpretation of the term to the respondents.³ We included primarily closed-form questions as well as a few open-ended questions for free responses. We also included standard demographic questions such as age, gender, education, income, region, and urbanicity. The final instrument included several dozen questions on a range of topics related to artificial intelligence (for more information, see the Appendix).

After we completed the instrument in English, we engaged cApStAn, a linguistic quality assurance agency with expertise in survey translation. We made several improvements based on their insights to minimize terminology that would be difficult to translate. In consultation with cApStAn, we also developed a translation style guide to ensure consistency and address complexities for particular concepts and/or languages. Our market research partner's in-country translation teams and/or third party vendors then translated the full instrument to all target languages while referring to the style guide. See Table 3 for the languages we offered. After the survey was complete, the responses were provided to the coding team to be coded in-language as described below. Illustrative quotes in this paper are verbatim (in the case of English language responses) or were prepared or reviewed by professional translators (in the case of non-English language responses).

3.2 Deployment

We selected a range of countries with different characteristics, such as stage of technological development, nature of the workforce, and varied development indices. The survey was fielded to online panels (groups of respondents who have agreed to participate in surveys over a period of time) representative of the online population in each country. Consistent with the best panels available for online market research, such panels tend to be broadly representative of the general population in countries with high access to technology, but less representative of the general population in countries with more limited access to technology; for example, in developing countries they tend to skew urban. Respondents were recruited using stratified sampling (a method of recruiting specific numbers of participants within demographic subgroups), with hard quotas on age and gender in each country. A summary of countries and demographics is provided in Table 3.^{4,5}

The median survey length was 21.4 minutes across all completions, including those who said they had never heard of AI in an early screening question and received a much shorter version of the survey. All respondents

²For logistical reasons we split data collection into two rounds: July 2019 (Australia, Canada, US, South Korea, France, Brazil, India, and Nigeria) and December 2019 (Germany, Finland, Singapore, Belgium, Spain, Poland, and China). Based on our experience with similar surveys and our knowledge of world events at the time, we do not expect the time interval between the two rounds had a substantial impact on the results.

³We note that in our pilot, we had two versions of the survey (one that defined AI and one that did not) and responses to subsequent questions were similar regardless of whether a definition had been provided.

⁴For compact layout, in all tables we use standard two-letter country codes, which are shown with full country names in Figure 1.

⁵The alert reader may notice the gender differences in India and Nigeria. Percentages were chosen to match benchmarks of the gender distribution of the online population in each country.

<i>Country</i>	DE	AU	FI	SG	BE	CA	US	KR
HDI Rank	6 th	8 th	11 th	11 th	14 th	16 th	17 th	23 rd
Languages offered	German	English	Finnish	Chinese, English	Dutch, French	English, French	English	Korean
Weighting	age, gender, education, region	age, gender, education, region	age, gender, education, region	age, gender, education	age, gender, education, region	age, gender, education, region	age, gender, education, region, race	age, gender, education, region
Respondents (<i>n</i>) All	1002	1000	1002	1000	1000	1500	1501	1000
Gender	51% men 49% women	49% men 51% women	48% men 52% women	54% men 47% women	54% men 46% women	47% men 53% women	49% men 51% women	52% men 48% women
Age, avg.	44	43	41	40	41	44	44	39
Age, stddev	15.4	15.3	15.2	12.7	15.1	16.0	17.5	12.4

<i>Country</i>	ES	FR	PL	BR	CN	IN	NG
HDI Rank	25 th	26 th	35 th	84 th	85 th	131 st	161 st
Languages offered	Spanish	French	Polish	Brazilian Portuguese	Chinese	English, Hindi	English
Weighting	age, gender, education, region	age, gender, education, region	age, gender, education, region	age, gender, education, region	age, gender, education, region	age, gender, education	age, gender, education
Respondents (<i>n</i>) All	1002	1001	1000	1503	1003	1500	1000
Gender	52% men 48% women	50% men 50% women	50% men 50% women	49% men 51% women	53% men 47% women	70% men 30% women	63% men 37% women
Age, avg.	42	43	41	34	38	30	31
Age, stddev	12.9	15.4	14.1	12.3	12.0	8.9	9.0

Table 3: Country details, respondent summary and demographics. All numbers unweighted.

received incentives in a point system or cash at an industry-standard amount for their market.

3.3 Data Processing and Analysis

3.3.1 Quality Checks

The market research firm conducted quantitative and qualitative checks to remove low quality responses on an ongoing basis until the quota was reached in each country. Example grounds for removal included being identified as a bot, speeding (answering substantially more quickly than the median time), or providing nonsensical or profane responses to open-ended questions.

Overall we removed 6.1% of responses for quality. After data collection was complete, standard procedures were followed to apply a modest weighting adjustment to each respondent so that the samples in each country are

more representative [5]. This weighting is reflected in the data shared in Section 4. The variables considered in weighting appear in Table 3.

3.3.2 Research Objective and Data

In this paper we focus on the following research objective: What sentiment do respondents have towards AI? Specifically, we present emergent themes, descriptive statistics, and illustrative quotes for the following open-ended question about sentiment:

‘What feelings or emotions come to mind when you hear the phrase Artificial Intelligence (AI)?’

When we present illustrative quotes, we also draw on responses from three additional open-ended questions (about description of AI, examples of AI, and any uncomfortable experiences with AI), as relevant responses and similar coding often applied across all of the open-ended questions.⁶

3.3.3 Coding and Analysis of Open-Ended Responses

We reviewed the open-ended responses from the pilot to identify emergent themes [4] and develop an initial codebook for all questions, then iterated as we reviewed responses from all countries to refine it as necessary. The open-ended responses were coded by our market research partner’s dedicated coding team or one of their third party coding vendors. The coding was done in the source language, with the exception of Dutch and Finnish which were coded based on English translations. As described in McDonald et al., a variety of different approaches may be employed to improve the reliability of qualitative analysis [43]. In our case, following best practices in public opinion research for coding against multiple languages, we used professional coders, followed an iterative process to continuously improve the codes, and performed a series of hierarchical quality checks. While coders were specialized by language, they worked together to ensure consistency, sharing notes in specialized coding software. Both we and our market research partner performed multiple levels of quality checks on the resulting coding, randomly sampling from all responses in each country as well as checking all instances of select codes.

We used an inductive approach to explore emerging themes and common patterns in the data [32]. For the open-ended question regarding the feelings or emotions the respondent associated with AI, we began by following the process described above; the resulting codebook for this question encompassed 92 codes (e.g. ‘Useful,’ ‘Skeptical,’ ‘AI takes over’) and specified that multiple codes could be assigned per response. After these codes were assigned and we reviewed the open-ended verbatim responses in detail, four thematic groups of codes emerged from the data as common and semantically distinct: **Exciting**, **Useful**, **Worrying**, and **Futuristic**. For example, the **Useful** group encompassed codes such as ‘Useful,’ ‘Helpful,’ ‘Productivity,’ etc. We assigned each of the 92 codes to exactly one of these four sentiment groups or Other accordingly. Other encompassed answers that were inarticulate, classified as unable to be coded, mentions of technology without any sentiment (e.g. “computer” or “technology”), and a long tail of other opinions on AI (for example “curiosity” or “surprise”). Based on the codes that each response had been assigned, each response was considered to be part of those group(s) – for example, if a response had been assigned the code ‘Helpful’ and the code ‘Concern,’ that response was part of the sentiment groups **Useful** and **Worrying**. A response that received *only* codes labeled Other appears in **None**.

3.3.4 Human Development Index

As the impact and use of AI expands worldwide, how people learn about, interact with, and use AI varies. People from developed countries (i.e. countries that are more industrialized and have higher per capita incomes, for

⁶The first question about sentiment was shown to all respondents, while the remaining three questions were only shown to respondents who reported that they had heard of AI before the survey.

example, Germany, Australia, Finland, or Singapore) have different circumstances than people from developing countries (i.e. countries that are less industrialized and have lower per capita incomes, for example, Brazil, China, India, and Nigeria), and this shapes how AI is perceived, adopted, and normalized globally [61, 56]. Therefore, we anticipated that there might be meaningful differences in AI perceptions associated with development level. We include the Human Development Index (HDI) Rank in Table 3.⁷

3.4 Limitations

We note several limitations of our methodology that should be considered when interpreting this work. First, it carries with it the standard issues attendant with survey methodology, such as the risk of respondents misunderstanding questions, poor quality translation, or respondents satisficing [33] or plagiarizing open-ended responses. We have worked to minimize these risks through piloting, use of open-ended questions in conjunction with closed-form questions, use of a translation style guide and translation review, and data quality checks. We also note that panels in India are well-known in the industry to be disproportionately likely to have a social desirability response bias (as defined in [33]), so optimism in the responses from India should be considered in that context. Second, online panels are not representative of the general population. While we have used a high standard of currently available online panels, we caveat our findings as not representative of the general population, particularly in Brazil, China, India, and Nigeria. Third, while members of the research team and/or market research partner team have experience conducting research in all markets studied, members of the team reside in Western countries. We have worked to minimize the risk of misinterpretation by collaboration and discussion with in-country partner teams but recognize that our interpretations may lack context or nuance that would have been more readily available to local residents.

4 Results

In this section we describe the sentiment groups that emerged from our analysis and present data on the frequency of their occurrence. Responses to the open-ended sentiment question were assigned to groups as described in Section 3.3.3. Many responses were brief and were assigned only one code, for example, responses such as “exciting” or “robot” would be assigned **Exciting** or **Futuristic**, respectively. However, responses were often more lengthy and received multiple codes. For example, a response such as “fear and excited at the same time” (US respondent) would be included in **Worrying** and **Exciting**, but not **Useful** or **Futuristic**.

Figure 1 and Table 4 present the prevalence of sentiment groups in each country. We now discuss each sentiment in turn.

4.1 Exciting

Responses in this group contained positive feelings about AI and often exhibited broad excitement or enthusiasm. These feelings were often direct statements of excitement, but also included other positive feelings such as joy or a sense of feeling blessed to have AI. **Exciting** sentiment was often associated with a sense of newness or expectation of substantial change. Sometimes respondents expressed excitement about improvements in daily life, and sometimes they anticipated broad improvement for humanity.

Excited to see where this tech goes in future, hope to see AI assist with everyday life in the home and in work –*Australia*⁸

⁷We show HDI ranks from the 2020 Human Development Report <http://hdr.undp.org/en/2020-report>, which uses HDI values from 2019, aligning with the dates of our survey deployment.

⁸Throughout the paper, we share complete verbatim responses (in some cases translated) and do not correct typographic or grammatical errors. The only exception is the quote in the title, which is a verbatim excerpt from a response that is shared in full in Section 4.6.

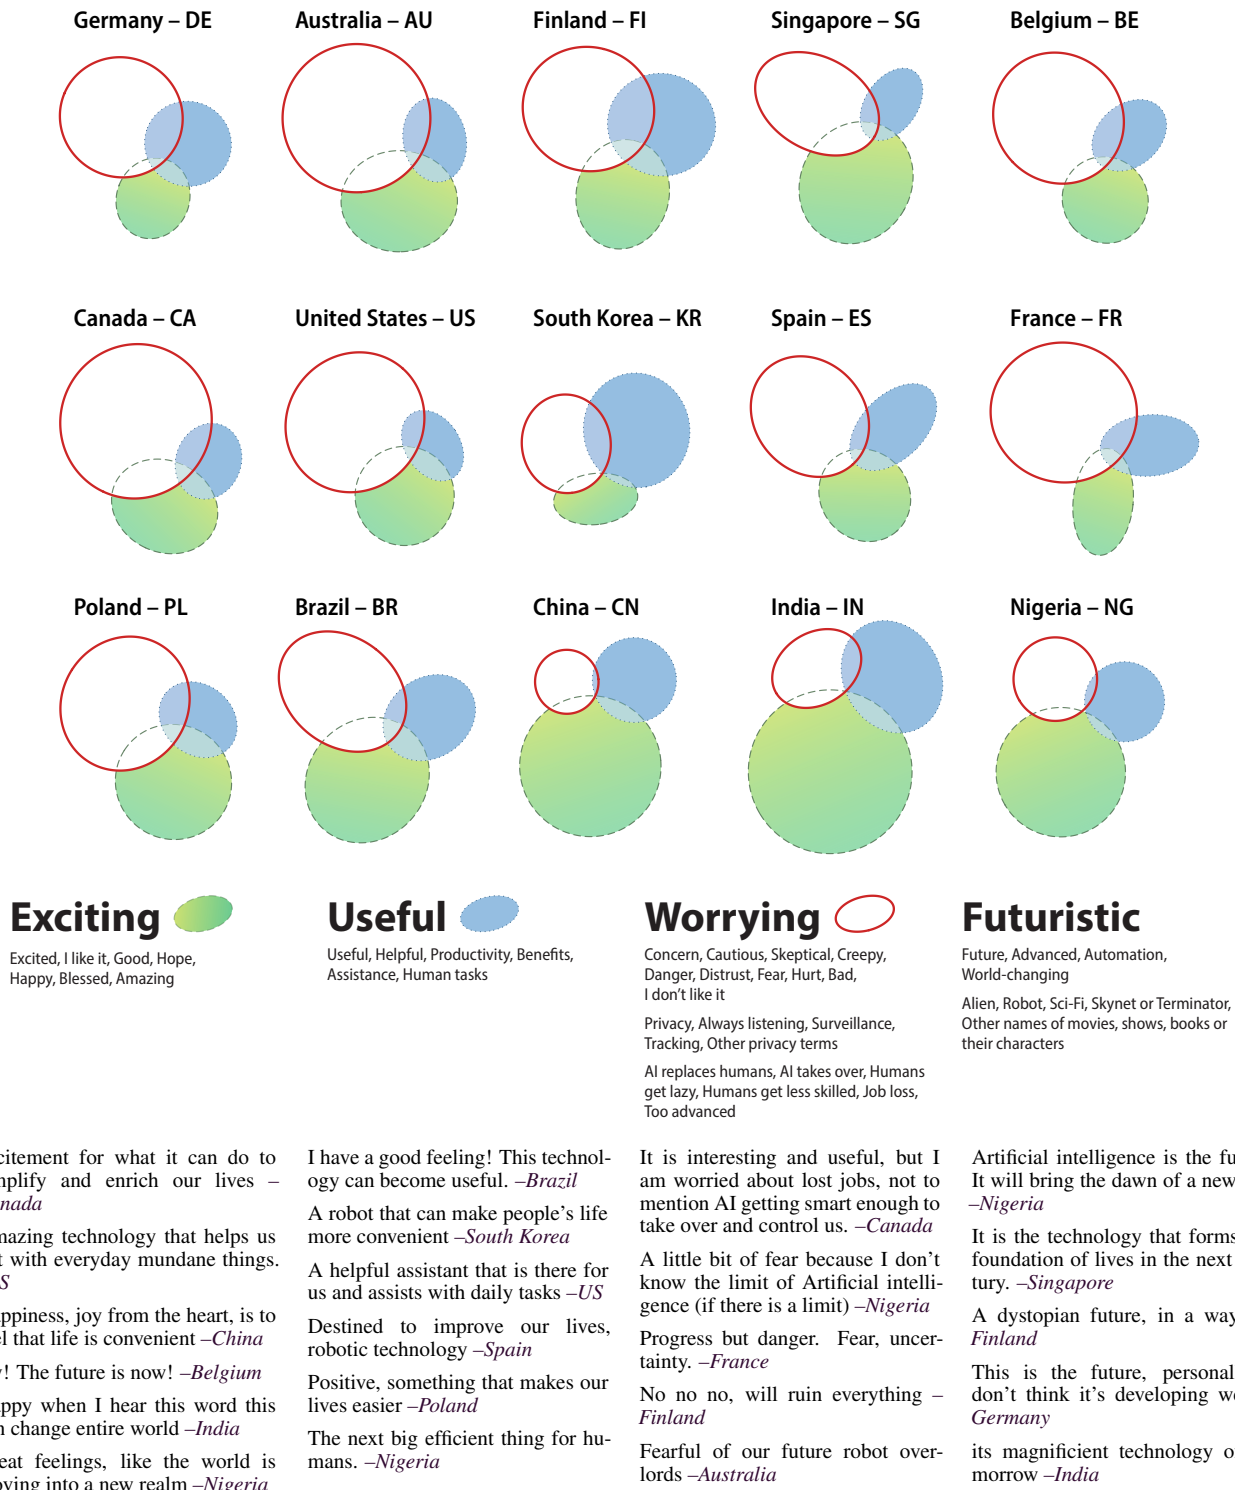


Figure 1: Description of our four sentiment groups, with the complete list of codes that comprises each, and example responses. While we use the responses to illustrate a particular sentiment, some of them fall in multiple sentiment groups, as sometimes occurred in our data set. At the top of the figure, we represent the overlap between the groups with Venn diagrams, using 3-Venn diagrams which exclude Futuristic for readability. The alert reader may wonder why we use oblong circles; these more accurately represent the area in the overlap. We use the method described in [44]. Countries are ordered by HDI.

Country	DE	AU	FI	SG	BE	CA	US	KR	ES	FR	PL	BR	CN	IN	NG
Respondents (n) All	1002	1000	1002	1000	1000	1500	1501	1000	1002	1001	1000	1503	1003	1500	1000
Exciting	9%	17%	14%	22%	12%	14%	15%	6%	14%	10%	25%	23%	28%	36%	25%
Useful	12%	9%	17%	7%	9%	9%	7%	19%	13%	11%	13%	14%	13%	17%	11%
Worrying	23%	31%	25%	20%	27%	33%	30%	14%	23%	31%	30%	21%	7%	9%	11%
Futuristic	17%	22%	24%	21%	15%	21%	19%	38%	28%	20%	22%	34%	24%	24%	19%
None	48%	38%	36%	38%	48%	39%	42%	31%	34%	39%	32%	25%	37%	27%	41%

Table 4: Weighted percentage of respondents from each country whose open-ended sentiment was coded to be in one of our groups. Respondents can appear in multiple sentiment groups. A respondent whose answers received only codes not in these groups appears in **None**.

Happiness, because it means humanity moves forward. –*Spain*

A bit of excitement, fascination, curiosity, but also somewhere deep a feeling of uncertainty, a thrill related to the effect that it may have on my future, however, that is mostly due to the influence of film rather than a result of conscious assessment of the benefits that AI will bring. –*Poland*

Some respondents were also excited about potential economic advantages for their country, and a few mentioned personal career opportunities that AI might provide.

4.2 Useful

Responses in this group expressed the belief that AI will be helpful and assist humans in completing tasks. **Useful** sentiment was generally associated with practical implications of AI. For example, respondents spoke of AI improving productivity in industrial settings. They also spoke of AI providing personal convenience, making people’s lives easier and more comfortable, assisting with daily life, enabling smart home technology, and helping people perform mundane tasks.

It is a kind of high technology that brings great convenience to our lives. –*China*

A tool of the future to make everyday life easier –*Finland*

Replaces man in thankless tasks –*Belgium*

Respondents also spoke of AI helping humanity by addressing large societal issues such as healthcare or the environment.

Really interesting. Hopefully it can solve energy issues and other large problems. –*Finland*

Progress, I know it will impact positively especially in the areas of health care. –*Nigeria*

4.3 Worrying

Many respondents shared that AI is **Worrying**, causing them concern, fear, or anxiety. Unlike the previous two groups, which each capture a relatively tight set of responses to AI in our open-ended data, this group comprises a wide range of negative emotional responses.

Do we need that? It scares me! –*Germany*

A dangerous game. –*Finland*

It’s progress, but I am not sure that it is so positive for society –*Belgium*

While some respondents spoke of general concerns, many spoke to specific aspects of AI they found **Worrying**. For example, some respondents were concerned that AI might challenge humans and take over society. Sometimes they suggested that popular culture had caused this concern. Correspondingly, they sometimes also spoke to the need for humans to control AI.

A threat to the future of humanity. –*Finland*

regret, it will make work and then the world disappear –*Belgium*

Can improve or end our lives. –*Germany*

A mixture of knowledge and fear. I know that it will help or is already helping in several important areas, but there is always that fear that one of these AIs will become too autonomous and turn against us. –*Brazil*

That was always coming. But because of all those films on TV with the AI, I still have my reservations. I mean you never know, right? –*Belgium*

Does no one watch movies, read, or anything to do with science fiction!!! It ALWAYS ends badly... there is just no good outcome, that I can see (for now at least), to an actual, fully fledged, AI. –*US*

Respondents also saw privacy concerns as a likely downside of AI. Sometimes they mentioned privacy in broad terms, but they often raised specific concerns such as worrying that products constantly listen to them, or concerns about being surveilled in the workplace.

A new frontier. Very exciting and scary at the same time. Lots to gain but will personal privacy be the price? –*Australia*

Could lead to total surveillance –*Germany*

A trending mobile app that undresses people. It violates privacy rules –*Nigeria*

IT DICTATED MY WEIGHT AND HEIGHT IN PUBLIC. –*India*

Installing a monitoring system in the office makes people very uncomfortable –*China*

Ads that show up on computers after visiting websites is one thing, but ads that show up after just talking about something makes me think my phone is listening in on my conversations –*US*

The phone's microphone recognizes speech and this information is used in marketing. Should I dare speak about sensitive matters near the phone at all –*Finland*

Respondents expected that AI would negatively impact the number of jobs available in the future. They perceived that AI may replace humans or make them less necessary in the workforce, and particularly associated robots with job loss due to their ability to perform human tasks. In rare cases, respondents shared personal experiences with automation-related job loss.

I feel that it has taken away jobs –*US*

A highly computerised potentially dangerous job stealing system of machinery operation –*Australia*

New technologies. Convenience in life. Reduction in jobs. –*South Korea*

Am happy about it but am still sceptical about it. This is because it might probably put some persons out of work –*Nigeria*

Unemployment comes to my mind when I hear the phrase Artificial Intelligence(AI). –*India*

Respondents also expressed concern that humans will become over-reliant on AI and become lazy, or that AI will minimize human contact and negatively impact personal relationships in the future.

This is a futuristic innovation that can help people but also make them too lazy –*Nigeria*

fear that during my lifetime I will be interacting more with AI than live humans –*US*

It helps the future by making things easier, but diminishes employment and human contact. –*France*

Country	DE	AU	FI	SG	BE	CA	US	KR	ES	FR	PL	BR	CN	IN	NG
Respondents (n) All	1002	1000	1002	1000	1000	1500	1501	1000	1002	1001	1000	1503	1003	1500	1000
Neutral	31%	24%	9%	24%	28%	26%	27%	14%	25%	25%	21%	11%	28%	9%	9%
I don't know	16%	17%	19%	5%	19%	15%	9%	8%	15%	17%	8%	11%	1%	4%	3%
Other	12%	7%	10%	6%	9%	7%	10%	14%	12%	9%	11%	9%	9%	15%	19%
Inarticulate	15%	5%	12%	10%	12%	7%	6%	4%	6%	19%	8%	4%	7%	7%	4%
Curious	3%	9%	8%	6%	5%	7%	6%	2%	10%	5%	22%	16%	16%	5%	7%
Intelligence	1%	4%	5%	9%	6%	3%	7%	10%	5%	2%	3%	9%	7%	13%	18%
Technology	3%	4%	6%	2%	2%	6%	5%	5%	7%	3%	3%	12%	11%	11%	10%
Fake	1%	3%	4%	3%	4%	4%	7%	3%	2%	1%	4%	2%	0%	6%	17%
Computer	3%	4%	5%	6%	3%	4%	7%	4%	2%	3%	2%	5%	1%	5%	6%
Device	1%	3%	4%	4%	1%	3%	2%	12%	3%	2%	3%	3%	1%	9%	6%

Table 5: The 10 most common codes within the **None** group, which includes all respondents who were assigned none of the codes in our four sentiment groups. Weighted percentages indicate how many respondents within the **None** group were assigned a given code, by country.

4.4 Future

Although **Futuristic** may not traditionally be seen as a sentiment, when we asked respondents to describe their feelings or emotions about AI, they organically responded at very high rates and it was clearly a strong association (15% to 38% across the countries). Responses in this group are not necessarily positive or negative towards AI, but rather are included for any mention of the futuristic nature of AI, whether by simply describing AI as advanced; mentioning robots, aliens, or other science-fiction concepts; or by referencing the future directly.

Some respondents who spoke of the future expressed that AI will be transformative, for example saying that it will usher in a new era and profoundly change society.

Something new. Something that will change the world –*Poland*

AI can change every aspect of human life. –*Singapore*

we are entering a new era. Very modern –*Canada*

AI will revolutionise the way we live in our future. –*India*

The expected future effects of AI were sometimes described as **Exciting**, **Useful**, **Worrying**, or some combination of these. We discuss mixed feelings further below.

A thing of the future that is sure to be of great use! –*Finland*

Better future –*Spain*

A big problem for humanity in the future. –*Poland*

Machines taking over humans!! :) on a serious note, A.I. is making things possible we thought were not possible a few years ago. Computers recognise faces and fingerprints of humans. Machines carry out so many things to assist humans. Everywhere we look there are examples of artificial intelligence around us. –*Australia*

AI is the new trend for technology, I myself being a tech geek i know that AI is soon going to change the whole world with it's endless possibilities. AI is the future of Mankind –*India*

4.5 None

The groups above do not cover all responses. Some responses were assigned only the 49 codes that we did not include in our sentiment groups, in which case they fell into the **None** group described in Section 3.3.3. These

included, for example, broad mentions of technology (responses such as “technology” or “computer”). See Table 5.

4.6 Mixed Feelings

As seen above, a given response sometimes contained multiple sentiments toward AI. In some responses, these were all positive sentiments, for example, excitement that AI would be helpful in daily life. However, a number of responses were more ambivalent. Sometimes such responses contrasted specific positive expectations (e.g. personal convenience or improved healthcare) with specific negative expectations (e.g. reduced privacy or job loss). Additionally, these mixed emotions were sometimes coupled with a sense of resignation or inevitability.

It is a wonderful and terrifying concept that is inevitable. –*Australia*

The future of our world in a way that represents both progress and destruction –*Canada*

optimistic that it will enhance peoples lives and bring about breakthroughs in many fields but also skeptical that people will lose their jobs and there will be an invasion of privacy –*Canada*

Life will be much more enjoyable, but I fear that we’d lost what makes us human. Robots will replace humans in various fields, but there are positive sides as well, a pet robot being one of them. –*South Korea*

I think people are afraid of it and it is the future –*Spain*

Artificial intelligence is something most people will come to depend on in a few decades. It will make life easier at the same time make people lose their jobs. But one I’m certain of is that AI is here to stay for good. –*Nigeria*

Some respondents suggested that the eventual impact of AI is not yet determined, and that multiple outcomes are possible.

Both a threat and an opportunity at the same time. –*Finland*

Mixture of amazement at the potential of this technology and concern about possible pitfalls. Could be the start of something amazing or the beginning of the end (a la Terminator). –*Australia*

Unsure about the net value - has lots of positives but also there are some very legitimate concerns. –*Canada*

It’s exciting to think about the things that could come about with AI that would make our lives easier and safer, but also scary of course, who knows how it will truly effect society –*US*

Some respondents also indicated that the effects of AI depend on how it is used, as well as who is using it.

Artificial intelligence worries me a bit because if it’s not used well it can be dangerous, it has no conscience or ethics, but I acknowledge that it is an amazing tool. –*France*

A bit excited because it makes job quite easy but again its scary if it the technology goes wrong like someone using it for evil purposes. –*Nigeria*

Artificial Intelligence is very useful for whole human world. But don’t use it in a bad way –*India*

Some respondents also spoke directly to responsible development of technology. For example, they emphasized the need to think about potential impacts of technology prior to development, or the need for regulation or ethical evaluation.

Angry that future concerns or negative impacts aren’t ever considered before technology is developed –*Australia*

It’s a positive thing, but it needs to be regulated. –*Belgium*

Unstoppable, but it requires technology ethics –*China*

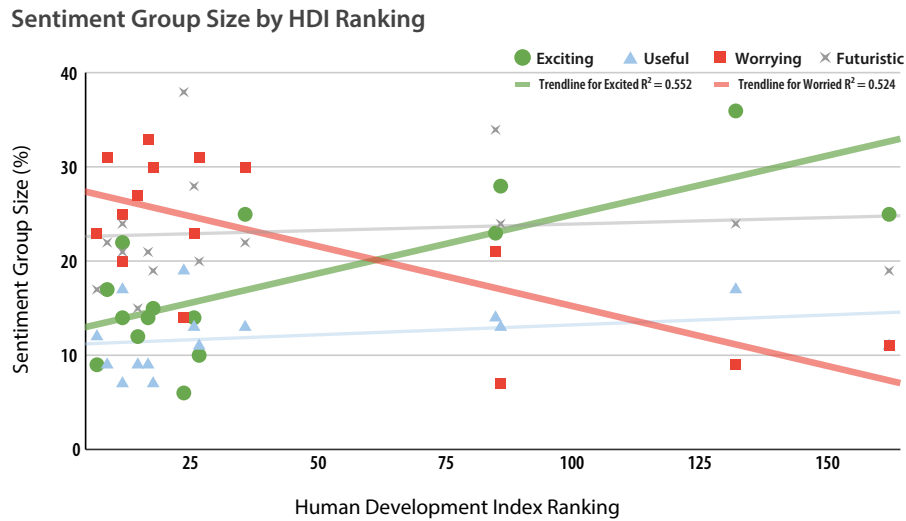


Figure 2: Scatter plot of weighted sentiment group size for each country, by HDI rank of the country. Trendlines shown for all four sentiment groups, with **Exciting** at $R^2 = 0.552$ and **Worrying** at $R^2 = 0.524$.

4.7 Country-Level Observations

We now turn to country-level observations, where we see strikingly different national patterns in response towards AI across the fifteen countries we studied. We visually represent the character of these differences in Figure 1.

Consistent with our expectation that developed countries (those most-developed, by HDI rank) would share similarities, the dominant sentiment groups in Germany, Australia, Finland, Belgium, Canada, the US, and France were **Worrying** followed by **Futuristic** (see Table 4). Spain had the same two dominant sentiments, although with **Futuristic** followed by **Worrying**. This resonates with claims that popular press and media narratives in Western, English-speaking regions have emphasized potential threats of AI [12, 25, 34].

By contrast, we see respondents in developing countries tend to take a more optimistic view of AI’s future effects. Respondents in China, India, and Nigeria were least likely to describe AI as **Worrying** and more likely to describe it as **Exciting**.

Singapore, Poland, and Brazil followed a different pattern, with more balanced numbers of **Worrying** and **Exciting**.

We can see this relationship more directly in Figure 2 where **Exciting** and **Worrying** show clear trends with HDI rank. **Futuristic** and **Useful** however do not seem to have a relationship with HDI, highlighting that the development of a country is just one factor in how public opinion towards AI is shaped.

South Korea has a unique profile among the countries surveyed, having the largest percentage in the both the **Useful** (19%) and **Futuristic** (38%) sentiment groups. South Korean respondents also had the lowest percentage of **Exciting** (6% versus 9-36% in all other countries). These findings are consistent with South Koreans’ high level of exposure to technology: South Korea boasts the world’s highest robot density [36], is one of the largest global investors in smart buildings [21], and may be “at the vanguard of a revolution in AI and big data healthcare” [46]. Consistent with this, South Korean respondents often mentioned AI assistants and home automation, which may contextualize AI as a more familiar, everyday technology:

AI is everywhere from hospitals to homes and cars. –*South Korea*

Use big data to make daily life more convenient. –*South Korea*

With just the smartphone, I can check the gas, temperature, and the foods in the fridge. –*South Korea*

Self-driving car, automated production, convenient daily life –*South Korea*

5 Discussion and Outlook

We conducted a large-scale investigation of sentiment towards AI across a range of countries. Rather than presupposing particular sentiment, we began with open-ended responses and looked for emergent themes. Our findings revealed sentiment groups as a distinguishing feature, with respondents in different countries finding AI to be **Exciting**, **Useful**, **Worrying**, and **Futuristic** to varying degrees. These groups provide one nuanced alternative to understanding people's feelings towards AI, rather than considering their orientation to AI as simply positive or negative. While some of these themes have been seen in other literature, here we have documented them occurring unprompted in 15 countries and added richer detail about the sentiment and the mechanisms which inspire this sentiment.

The spontaneously generated open-ended responses reflect a number of key dialogues that have appeared in public discussions and the media [12, 18, 59, 51], for example, that AI offers significant improvements for health; that AI is associated with privacy issues, job loss, and social isolation; and that AI could be either a significant boon or a significant threat to humanity. The data provide some indication of the ways in which these concepts, as well as different sources of information (e.g. fiction, news reports, or personal experience) influence sentiment.

This suggests many fruitful avenues for further exploration. For example, it would be valuable to more formally measure and analyze the relationship between media and pop culture narratives in different countries and the presence of these sentiment groups, as well as tracing the relationship and movement of narratives across countries. Further, it would be useful to explore other factors that likely influence these sentiment groups, such as country culture and economy; institutional trust [17, 67]; presence, awareness, and availability of AI technologies such as customer service chatbots, personal assistants, and more; and personal, formative experiences using AI technology. It would also be worthwhile to explore how sentiment groups affect behavior such as adoption of AI technologies and public opinion on topics such as research funding and regulation.

Public opinion has the potential to shape (and be shaped by) technology development processes and decisions. For example, public opinion can affect whether the public supports research funding for AI. As another example, a negative opinion of a particular technology may discourage consumers from purchasing it. Conversely, new product offerings that rely on AI may influence the public's opinion of AI.

While public opinion can be a beneficial influence, it has also been argued that in some cases it can have suboptimal effects. For example, public misperception or unrealistic expectations of AI may lead to unfounded fears or disappointment, resulting in unwarranted rejection of technology or a lack of support for public funding [6, 12]. But how might one characterize the "legitimacy" of public opinion, and to what extent is such characterization a meaningful endeavor? Many issues related to AI are complex questions on which even experts can disagree. And even if experts are in alignment with each other, but not with public opinion, the public may be considering perspectives or values not taken into account by experts [67]. It is therefore a complex question how best to interpret or engage with public opinion on a given issue, or whether it might be helpful to influence it. In some cases it may be beneficial to provide the public additional information, while in others it may be more beneficial for researchers and developers to shift the perspectives and values driving the development of AI.

Possible interventions might include educational efforts in areas in which the public may benefit from additional information. Beyond that, however, our findings align with calls to develop technology that supports public values. For example, many respondents were concerned about negative impacts of AI on privacy, reinforcing the value of continued emphasis on designing and developing AI with privacy in mind, concordant with discussion of privacy by design in the EU General Data Protection Regulation (GDPR).⁹ The privacy discussion continues to evolve quickly, and best practices for AI technologies continue to be actively explored in the academic, legal, and policy communities, offering many opportunities for advances in this area. Further, our findings also suggest ways in which the design and development of particular technologies may have a favorable impact on public

⁹<https://eugdpr.org/>

opinion. For example, our findings point to the value of emphasizing AI’s application to healthcare in product and research investments as well as communications. As another example, future research could explore the conditions facilitating South Korea’s unusually strong impression of AI as **Useful**, to gain insight into whether or how this sentiment might resonate elsewhere via communications or technological offerings.

Acknowledgments

We thank Dan Altman, Elie Bursztein, Ed Chi, Charina Chou, Jen Gennai, Reena Jana, Jake Lucchi, Ken Rubinstein, and Kurt Thomas for valuable contributions to this work. We also thank the cApStAn team for their important contributions to linguistic quality and the Ipsos team for their excellent work fielding the survey.

References

- [1] Christophe Abrisart, Yoshua Bengio, Guillaume Chicoisne, Nathalie de Marcellis-Warin, Marc-Antoine Dilhac, Sebastien Gambs, Vincent Gautrais, Martin Gibert, Lyse Langlois, Francois Laviolette, Pascale Lehoux, Jocelyn Maclure, Marie Martel, Joelle Pineau, Peter Railton, Catherine Regis, Christine Tappolet, and Nathalie Voarino. Montreal declaration for responsible development of artificial intelligence, 2018.
- [2] ARM | Northstar. AI today, AI tomorrow. awareness and anticipation of AI: A global perspective, 2017.
- [3] Chelsea Barabas, Madars Virza, Karthik Dinakar, Joichi Ito, and Jonathan Zittrain. Interventions over predictions: Reframing the ethical debate for actuarial risk assessment. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 62–76, 2018.
- [4] Hugh Beyer and Karen Holtzblatt. *Contextual Design: Defining Customer-centered Systems*. Elsevier, 1997.
- [5] Paul P. Biemer and Sharon L. Christ. Weighting survey data. In Edith D. de Leeuw, Joop J. Hox, and Don A. Dillman, editors, *International Handbook of Survey Methodology*, pages 317–341. Lawrence Erlbaum Associates New York, NY, 2008.
- [6] Blumberg Capital. Artificial intelligence in 2019: Getting past the adoption tipping point, 2019.
- [7] Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company, 2014.
- [8] Taina Bucher. The algorithmic imaginary: exploring the ordinary affects of Facebook algorithms. *Information, Communication & Society*, 20(1):30–44, 2017.
- [9] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 77–91, 2018.
- [10] Daniel Castro. The U.S. may lose the AI race because of an unchecked techno-panic. *Center for Data Innovation*, March 2019.
- [11] Stephen Cave, Kate Coughlan, and Kanta Dihal. “Scary Robots”: Examining public responses to AI. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES 2019)*, pages 331–337, 2019.
- [12] Stephen Cave, Claire Craig, Kanta Sarasvati Dihal, Sarah Dillon, Jessica Montgomery, Beth Singler, and Lindsay Taylor. *Portrayals and perceptions of AI and why they matter*. The Royal Society, 2018.
- [13] Stephen Cave, Kanta Dihal, and Sarah Dillon. *AI Narratives: A History of Imaginative Thinking about Intelligent Machines*. Oxford Scholarship Online, 2020.
- [14] CBS News. 60 Minutes/Vanity Fair poll: Artificial intelligence, March 2016.
- [15] Center for Information Technology Policy and University Center for Human Values. Princeton dialogues on AI and ethics case studies. <https://aiethics.princeton.edu/case-studies>.
- [16] Stevie Chancellor, Michael L. Birnbaum, Eric D. Caine, Vincent M. B. Silenzio, and Munmun De Choudhury. A taxonomy of ethical tensions in inferring mental health states from social media. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, page 79–88, 2019.

- [17] Yi-Ning Chen and Chia-Ho Wen. Impacts of attitudes toward government and corporations on public trust in artificial intelligence. *Communication Studies*, 72(1):115–131, 2020.
- [18] Ching-Hua Chuan, Wan-Hsiu Tsai, and Su Cho. Framing artificial intelligence in American newspapers. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES 2019)*, pages 339–344, 2019.
- [19] Michael A. DeVito, Jeremy Birnholtz, and Jeffery T. Hancock. Platforms, people, and perception: Using affordances to understand self-presentation on social media. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 740–754, 2017.
- [20] Thomas G. Dietterich and Eric Horvitz. Rise of concerns about AI: Reflections and directions. *Communications of the ACM*, 58(10):38–40, 2015.
- [21] DuckerFrontier. Smart building trends in 2019: Part 1, April 2019.
- [22] Tim Dutton. An overview of national AI strategies. *Medium*, June 2018.
- [23] Edelman. 2019 Edelman AI survey, March 2019.
- [24] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. “I always assumed that I wasn’t really that close to [her]”: Reasoning about invisible algorithms in news feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 153–162, 2015.
- [25] Ethan Fast and Eric Horvitz. Long-term trends in the public perception of artificial intelligence. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [26] Jessica Fjeld, Nele Achten, Hannah Hilligoss, Adam Nagy, and Madhulika Srikumar. Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *Berkman Klein Center Research Publication*, 2020.
- [27] Markkula Center for Applied Ethics. Ethics in technology practice. <http://www.scu.edu/ethics-in-technology-practice/>.
- [28] Cary Funk, Alec Tyson, Brian Kennedy, and Courtney Johnson. Science and scientists held in high esteem across global publics. *Pew Research Center*, September 2020.
- [29] Colin Garvey and Chandler Maskal. Sentiment analysis of the news media on artificial intelligence does not support claims of negative bias against artificial intelligence. *OMICS: a Journal of Integrative Biology*, 2019.
- [30] Stephen Hawking, Max Tegmark, and Frank Wilczek. Transcendence looks at the implications of artificial intelligence - but are we taking AI seriously enough? *The Independent*, May 2014.
- [31] High-Level Expert Group on Artificial Intelligence. Ethics guidelines for trustworthy AI, 2019.
- [32] Timothy R. Hinkin. A brief tutorial on the development of measures for use in survey questionnaires. *Organizational Research Methods*, 1(1):104–121, 1998.
- [33] Allyson L. Holbrook, Melanie C. Green, and Jon A. Krosnick. Telephone versus face-to-face interviewing of national probability samples with long questionnaires: Comparisons of respondent satisficing and social desirability response bias. *Public Opinion Quarterly*, 67(1):79–125, 2003.
- [34] Eric Horvitz and Bart Selman. Interim report from the panel chairs: AAAI Presidential Panel on long-term AI futures. In *Singularity Hypotheses*, pages 301–308. Springer, 2012.
- [35] Institute for the Future and Omidyar Network. Ethical OS Toolkit. <https://ethicalos.org>.
- [36] International Federation of Robotics. Robot density rises globally, Feb 2018.
- [37] Ipsos. Widespread concern about artificial intelligence, 2019.
- [38] Anna Jobin, Marcello Ienca, and Effy Vayena. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1:389–399, September 2019.

- [39] Shivaram Kalyanakrishnan, Rahul Alex Panicker, Sarayu Natarajan, and Shreya Rao. Opportunities and challenges for artificial intelligence in India. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (AIES 2018)*, pages 164–170. ACM, 2018.
- [40] Patrick Gage Kelley, Yongwei Yang, Courtney Heldreth, Christopher Moessner, Aaron Sedley, Andreas Kramm, David T. Newman, and Allison Woodruff. Exciting, useful, worrying, futuristic: Public perception of artificial intelligence in 8 countries. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (AIES '21)*, page 627–637, 2021.
- [41] Lloyd’s Register Foundation. World Risk Poll report 2019, 2020.
- [42] Pamela McCorduck. *Machines Who Think*. A K Peters/CRC Press, 1979.
- [43] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. In *Proceedings of the 22nd ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW 2019)*, 2019.
- [44] Luana Micallef and Peter Rodgers. eulerAPE: drawing area-proportional 3-Venn diagrams using ellipses. *PloS one*, 9(7):e101717, 2014.
- [45] Mozilla. We asked people around the world how they feel about artificial intelligence. Here’s what we learned., 2019.
- [46] Nature. AI and big data healthcare in Korea. *Nature Focal Point*, March 2019.
- [47] Lisa-Maria Neudert, Aleksi Knuutila, and Philip N. Howard. Global attitudes towards AI, machine learning & automated decision making: Implications for involving artificial intelligence in public service and good governance. *Oxford Internet Institute*, 2020.
- [48] Nils J. Nilsson. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge University Press, 2009.
- [49] Northeastern University and Gallup. Optimism and anxiety: Views on the impact of artificial intelligence and higher education’s response, January 2018.
- [50] OpenAI. Better language models and their implications. *OpenAI Blog*, Feb 2019.
- [51] Leila Ouchchy, Allen Coin, and Veljko Dubljević. AI in the headlines: The portrayal of the ethical issues of artificial intelligence in the media. *AI & Society*, 35:927–936, 2020.
- [52] Sundar Pichai. AI at Google: our principles, June 2018. <http://www.blog.google/technology/ai/ai-principles/>.
- [53] Emilee Rader and Rebecca Gray. Understanding user beliefs about algorithmic curation in the Facebook news feed. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 173–182, 2015.
- [54] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, page 469–481, 2020.
- [55] Neil J. Salkind, editor. *Encyclopedia of Research Design*, volume 1. Sage, 2010.
- [56] Nithya Sambasivan and Jess Holbrook. Toward responsible AI for the next billion users. *Interactions*, 26(1):68–71, 2019.
- [57] Javier Sánchez-Monedero, Lina Dencik, and Lilian Edwards. What does it mean to ‘solve’ the problem of discrimination in hiring?: Social, technical and legal perspectives from the UK on automated hiring systems. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, page 458–468, 2020.
- [58] Paul Scharre. *Army of None: Autonomous Weapons and the Future of War*. WW Norton & Company, 2018.
- [59] Peter Stone, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, et al. Artificial intelligence and life in 2030. *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*, page 52, 2016.
- [60] The European Commission. Special Eurobarometer 460: Attitudes towards the impact of digitisation and automation on daily life, May 2017.
- [61] United Nations. World economic situation and prospects: Country classification, 2014.

- [62] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *Proceedings of the Eighth Symposium on Usable Privacy and Security (SOUPS) 2012*, 2012.
- [63] Jeffrey Warshaw, Nina Taft, and Allison Woodruff. Intuitions, analytics, and killing ants: Inference literacy of high school-educated adults in the US. In *Proceedings of the Twelfth Symposium on Usable Privacy and Security (SOUPS) 2016*, pages 271–285, 2016.
- [64] Darrell M. West. Brookings survey finds divided views on artificial intelligence for warfare, but support rises if adversaries are developing it. *Brookings*, August 2018.
- [65] Max L. Wilson, Ed H. Chi, Stuart Reeves, and David Coyle. RepliCHI: The Workshop II. In *CHI Extended Abstracts '14*, page 33–36, 2014.
- [66] Jia You. A 100-year study of artificial intelligence? Microsoft Research’s Eric Horvitz explains. *Science*, January 2015.
- [67] Baobao Zhang. Public opinion toward artificial intelligence. In Justin Bullock, Baobao Zhang, Yu-Che Chen, Johannes Himmelreich, Matthew Young, Antonin Korinek, and Valerie Hudson, editors, *Oxford Handbook on AI Governance*. Oxford University Press, 2022.
- [68] Baobao Zhang and Allan Dafoe. Artificial intelligence: American attitudes and trends. *Available at SSRN 3312874*, 2019.

Appendix: Select Questions

Note that some questions were modified from or replicate other questions in the literature or the canon of public opinion surveys. For additional select questions used in the instrument see arXiv:2001.00081

Unaided Sentiment

{Ask All}

What feelings or emotions come to mind when you hear the phrase Artificial Intelligence (AI)?

{Open-end}

Knowledge

{Ask All}

How much do you know about Artificial Intelligence (AI)?

- A lot
- A moderate amount
- A little
- Heard of AI, but know nothing about it
- Never heard of AI

Unaided Description

{Do NOT ask if “Never heard of AI” in Knowledge question}

In your own words, please describe Artificial Intelligence (AI).

{Open-end}

Unaided Examples

{Do NOT ask if “Never heard of AI” in Knowledge question}

Please list some examples of how Artificial Intelligence (AI) is used today.

{Open-end}

Uncomfortable Experience

{Do NOT ask if “Never heard of AI” in Knowledge question}

Have you ever had an experience with AI-related technology that made you feel uncomfortable?

- Yes

- No
- Not sure

Unaided Description of Uncomfortable Experience

{Ask if “Yes” to Uncomfortable Experience}

What happened, and what was the outcome? Please describe your experience with AI that made you feel uncomfortable.

{Open-end}