NOVEMBER 2020

"Cool Projects" or "Expanding the Efficiency of the Murderous American War Machine?"

AI Professionals' Views on Working With the Department of Defense

CSET Issue Brief

CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

AUTHORS

Catherine Aiken
Rebecca Kagan
Michael Page

Note: Title quotations are taken from survey responses. They do not reflect the views of the authors or CSET. "Cool projects" was a response to the question, "What reason do you find most convincing for working on a DOD AI contract or grant?" "Expanding the efficiency of the murderous American war machine" was a response to the question, "What reason do you find most convincing for not working on a DOD AI contract or grant?" The quotations are meant to illustrate two poles of thinking on this issue and do not represent the central findings of this report.

Executive Summary

In a world where computing technologies such as artificial intelligence are increasingly critical for government operations and national security, the nature of the relationship between the U.S. tech industry and the Department of Defense is important. Recent highly publicized incidents of tech employee pushback against DOD-funded AI projects have created the perception of a rift between the tech industry and DOD. If true, such a rift would have serious implications for public policy, national security, and the future of AI. But are recent examples of employee pushback representative of the feelings of the entire tech industry? Specifically, how do tech professionals with AI-relevant skills feel about working on DOD-funded projects?

To better understand the tech industry-DOD relationship, CSET surveyed 160 U.S. AI industry professionals with AI-relevant skills (AI professionals) about their views on working on AI projects funded by DOD grants or contracts (DOD-funded AI projects). We find that AI professionals hold a broad range of views about DOD generally and working on DOD-funded AI projects specifically.

Key findings from the survey include:

Most AI professionals are positive or neutral about working on DOD-funded AI projects. Nearly 40 percent are neutral and 38 percent are extremely or somewhat positive. Less than a quarter of respondents feel extremely or somewhat negative about working on a DOD-funded AI project.

AI professionals consider interesting research opportunities and the ability to do good to be the most compelling reasons to work on DOD-funded AI projects. Respondents who feel positively about working on DOD-funded AI projects also consider the ability to influence DOD to be a compelling benefit.

AI professionals' discomfort with how DOD will use the technology and concerns about causing harm were the most common reasons not to work on DOD-funded AI projects. Respondents who feel negatively about working on DOD-funded AI projects were particularly concerned about DOD's use of the technology.

AI professionals are more willing to work on DOD-funded AI projects with humanitarian applications, as opposed to battlefield or back-office applications. Additionally, the ability to provide a global

humanitarian benefit is the most frequently cited factor that would increase willingness to work on a DOD AI contract.

AI professionals' willingness to work on DOD-funded AI projects is also increased if the funding is used solely for basic research. Stronger commitments to safety and privacy as well as projects with defensive applications increase willingness for around 40 percent of AI professionals.

AI professionals who worked at an employer with DOD contracts and those more familiar with DOD tend to be more positive about working on DOD-funded AI projects. Just over half of respondents have worked or currently work at an employer with DOD contracts and about a quarter worked directly on a DOD-funded project. About half of surveyed professionals consider themselves at least somewhat familiar with DOD.

AI professionals consider their employer and academia to be the most trustworthy actors to develop AI in the public interest. A minority of AI professionals trust the U.S. government and military to develop AI in the interest of the public. Compared to surveys of the broader American public, AI professionals have lower levels of trust in the U.S. military's ability to develop AI but higher levels of general trust in the U.S. government.

Table of Contents

Introduction

Historically, DOD has been the primary funder and consumer of military-relevant technologies developed in the private sector. The relationship has been mutually beneficial: DOD funded speculative research and development in areas without obvious commercial value that ultimately bore both military and commercially relevant advances.

Today, the increasing role of computing technologies, such as AI, in national defense and the rise of the tech industry are fundamentally reshaping that dynamic. The tech industry leads the development of military-relevant AI technologies, for which DOD is no longer the primary funder or consumer. Unable to compete with the resources available in the private sector, DOD's ability to influence the direction of AI development and leverage new advances toward military ends is diminishing.

In response to these shifting dynamics, relationship building with the tech industry has become a priority for DOD. In 2015, DOD launched the Defense Innovation Unit, with offices in Silicon Valley, to connect DOD with technology companies and build a national security innovation base to accelerate DOD's adoption of commercial technology. The following year, DOD announced the Defense Innovation Board, chaired by former Google CEO Eric Schmidt, with a two-year, renewable mandate to bring the best of Silicon Valley to the U.S. military. In February 2020, DOD adopted five AI ethical principles based on recommendations by the Defense Innovation Board, suggesting continued prioritization of relationship building with the tech industry.[1]

Yet, recent DOD and tech industry partnerships have faced high-profile pushback. In 2017, DOD launched a pilot AI program, the Algorithmic Warfare Cross-Functional Team, known as Project Maven, with the goal of developing computer vision algorithms for object detection and classification.[2] Google was awarded an 18-month contract to work on the project. However, the company faced significant pushback in 2018 when thousands of employees signed a letter to Google CEO Sundar Pichai saying, "Google should not be in the business of war."[3] Employees staged a walkout, and several resigned in protest. Media coverage honed in on the employee backlash, describing it as a "culture clash between Silicon Valley and the federal government."[4] In response, Google announced it would not renew the Project Maven contract or pursue certain government contracts in the future, saying it "will not support the use of AI for weaponized systems."[5]

Google was not alone in facing employee pushback over DOD contracts. In 2018, Microsoft announced it would bid on DOD's Joint Enterprise Defense Infrastructure cloud project, known as Project JEDI.[6] In the announcement, Microsoft President Brad Smith articulated two key reasons for the company's engagement with DOD: to provide necessary support for those who defend the United States and to allow Microsoft to engage in debate about how new technologies can be used responsibly. Again, some employees pushed back, calling on Microsoft to drop the bid on Project JEDI.[7] Despite protests, Microsoft did not withdraw its bid, and DOD awarded it the contract in October 2019.

These high profile examples of employee protests against DOD contracts suggest mounting tension between the tech industry and DOD, and foster the narrative of AI professionals unwilling to work in collaboration with the U.S. military.[8] The media amplified the narrative of a divide, reporting, for example, that "the [tech industry-DOD] rift is real, deep, and a long time coming."[9] Dueling opinion pieces argued that Silicon Valley should either "go to war" or avoid collaboration with the Pentagon entirely.[10] Meanwhile, major tech and policy figures staked their claims in the debate. Defense Innovation Board chairman Eric Schmidt stated the United States needs "unprecedented partnerships between government and industry." Amazon founder and CEO Jeff Bezos agreed, stating, "if big tech is going to turn their backs on the Department of Defense, this country is in trouble."[11] As recently as August 2020, Palantir CEO Alexander Karp wrote "if we are going to ask someone to put themselves in harm's way, we believe that we have a duty to give them what they need to do their job."[12] In contrast, groups like the Electronic Frontier Foundation lauded Google's break from Project Maven as a "big win for ethical AI principles" and Canadian computer scientist and AI pioneer Yoshua Bengio spoke out "very firmly against" military use of AI technologies, expressing concern that military organizations "put duty before morality."[13]

Still others questioned the notion of a tech industry-DOD divide entirely, arguing that relations remain strong. When asked about the DOD relationship with Silicon Valley in 2018, Defense Innovation Unit's Navy lead Captain Sean Heritage stated, "this is really a non-story." Others in government echoed that response, calling concerns about the potential rift overblown.[14]

Growing interest in the relationship between the tech industry and DOD has led to some research exploring these arguments. A 2016 study by the Center for a New American Security and The Copia Institute interviewed and surveyed a small number of tech and security policy professionals about their

views of the state of the tech-security relationship. The researchers found evidence of a strained, adversarial relationship stemming from a lack of productive dialogue and mutual understanding.[15] However, a more recent survey of U.S. tech professionals found that a majority agreed that "tech companies should work with the US government on military projects."[16] Similarly, a recent analysis of U.S. federal contracting data from 2016-2020 found no systemic divide between Silicon Valley and Washington. It concluded that recent narratives decrying the divide are based on anecdotal evidence and that "tech companies, universities, and tech-oriented traditional defense contractors, by and large, appear to be willing and able to modernize US battle networks."[17]

As a whole, the research on the relationship between DOD and the tech industry is limited and, at times, contradictory. We are left with conflicting narratives and limited data. Media narratives and anecdotal evidence continue to underline the perception of a rift, but the true extent of a divide and underlying drivers of division remain unclear. Our research aims to fill this gap by assessing AI professionals' views toward working with DOD.

Survey Methodology

What are the views of AI professionals on working with DOD, and what factors account for variation in those views? To address these questions, we surveyed a sample of AI-skilled employees of U.S.-based AI companies (AI professionals) about their views on DOD.[18] We randomly sampled AI professionals working at small and large companies with the highest concentration of AI-skilled employees (AI companies) in the U.S. metropolitan areas with the highest number of AI-skilled workers (AI hubs).[19] We identified AI professionals through LinkedIn based on reported skills, employer, and location.[20] We collected email addresses for identified professionals through online profiles and email finder services. To better understand our target group and refine our survey questions, we also conducted two focus groups; the first in San Francisco in November 2019 and the second in Boston in December 2019.[21] We sent the survey by email to approximately 4,000 professionals between March and May 2020.

The survey asked AI professionals about their willingness to work on DOD-funded AI projects, the perceived benefits and downsides of working on such projects, views toward various government entities, and thoughts about the implications of their work. The survey also included an experiment that tested respondents' willingness to work on a hypothetical DOD-funded AI project that varied in its end-use application (department back-office, humanitarian relief, or battlefield enhancement) and project scope (U.S. or global), and asked what actions respondents would take in support of or against the project. Lastly, the survey asked for background information, including political interest and involvement, past experience working on DOD-funded projects, and familiarity with DOD-related topics. See Appendix B for the full survey questionnaire.

The survey was distributed over several pilot waves and a main distribution. To maximize response rates for our primary questions of interest, some respondents completed a shortened version of the survey. That version asked respondents their willingness to work on DOD-funded AI projects, the factors that would increase their willingness to work on such projects, and included the survey experiment. In total, 28 percent of respondents completed only the short version of the survey, 42 percent of respondents completed the extended survey, and 31 percent partially completed the extended survey. Responses from pilot distributions and partial responses are included in our analysis and the number of respondents for each question is listed in the figure captions.

We collected 160 responses, for a four percent response rate.[22] In an effort to understand the source of the low response rate, we tested a pilot survey invitation that varied the description of the survey topic and we asked nonrespondents to provide their reason(s) for not taking the survey. We found no discernible difference in response rate by removing references to national defense and security from the invitation.[23] We also found that no single reason dominated among the small number of nonrespondents who offered a reason for not taking the survey.[24] While not definitive, this suggests that low response propensity may be more about the population of interest than the survey topic. Despite the low response rate, our sample is representative of AI professionals working in current U.S. AI hubs in terms of professional activities, employer size, and work location (see Table A in Appendix). While our respondents are representative of AI professionals in terms of these observable characteristics, we cannot know if they are representative in terms of unobserved characteristics. It may be that our respondents and non-respondents differ attitudinally in ways relevant to their views on DOD. Because this may skew our results in unpredictable ways, we emphasize that our results are only a first step toward understanding industry perceptions and are not definitive conclusions.

Findings

Our findings call into question the narrative of widespread unwillingness on the part of AI professionals to work with DOD. We find that AI professionals hold a range of views about working with DOD. While some had strong negative views, more had positive views or were neutral regarding working on DOD-funded AI projects. By several measures, respondents are open to working on DOD-funded AI projects, motivated primarily by the content of the work, by interesting problems and access to unique resources. Yet, willingness is not unconditional and is matched with concerns about how DOD will use the technology. AI professionals have a preference for collaborating on basic research and humanitarian projects and against work that they perceive as causing harm. Overall, instead of finding a tech industry that is unified by its unwillingness to work with DOD or desire to 'go to war,' we find an industry diverse in its views toward collaborating with DOD.

AI professionals

A majority of respondents, 62 percent, work at large companies with more than 10,000 employees. In terms of primary professional activities, 60 percent of respondents selected engineering as a primary professional activity, 51 percent selected applied research, and 48 percent selected data analysis.[25] As shown in Table 1, San Francisco was the most common AI hub where respondents work, reflecting the current geographical breakdown of AI professionals working in AI hubs. A majority of respondents are between the ages of 25 and 44, and only 11 percent are over the age of 55. In terms of citizenship, 81 percent who report citizenship are U.S. citizens.[26]

Table 1. Respondent Work Location

AI Hub	Respondents	AI Professionals
San Francisco	32%	39%
New York City	19%	20%
Seattle	16%	13%
Boston	13%	10%
Washington, D.C.	13%	9%
Los Angeles	8%	9%
Total	160	3,995

Attitudes toward working on DOD-funded AI projects

To begin, we asked respondents to indicate their general feelings toward working on an AI project funded by a DOD grant or contract. Results are displayed in Figure 1. Most respondents were neutral or positive about working on a DOD-funded AI project, while about a quarter were negative. A small fraction of respondents felt extremely positive or extremely negative, but more were extremely positive than were extremely negative (see Figure A in Appendix). Respondents were especially positive if a project was specifically for basic research, as shown in Figure 2.[27]

Figure 1. AI Professionals' Views on Working on DOD-Funded AI Projects Vary

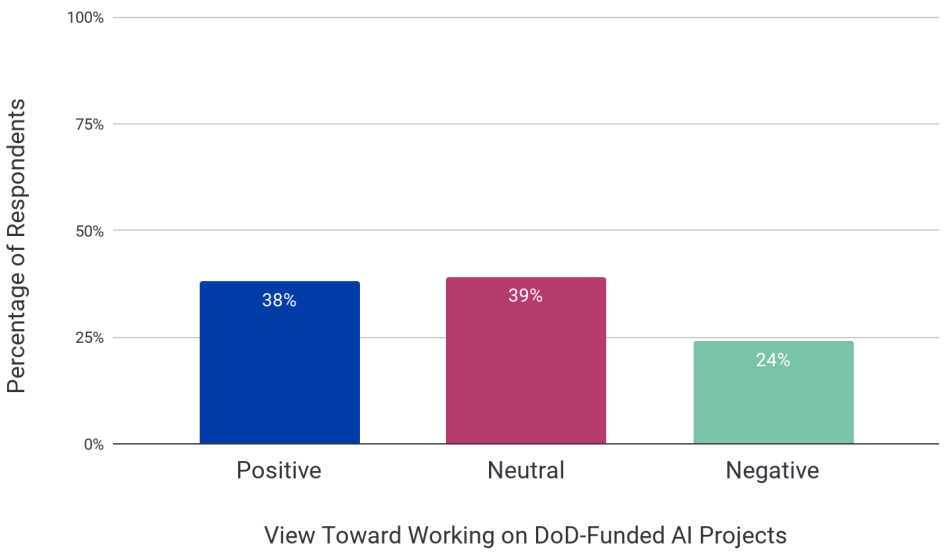

Question asked, "Do you generally feel negative, neutral, or positive about working on an AI project funded by a DOD grant or contract?" Figure combines responses from pilot version which offered a 3-point scale of "negative," "neutral," or "positive" and the main survey which offered a 5-point scale of "extremely negative," "somewhat negative," "neither negative nor positive," "so mewhat positive," or "extremely positive." Respondents: 140. Proportions do not sum to 100 percent due to rounding. Source: CSET AI-DOD Survey 2020.

Figure 2. AI Professionals are More Positive About DOD Funding for Basic Research Compared to Applied Research

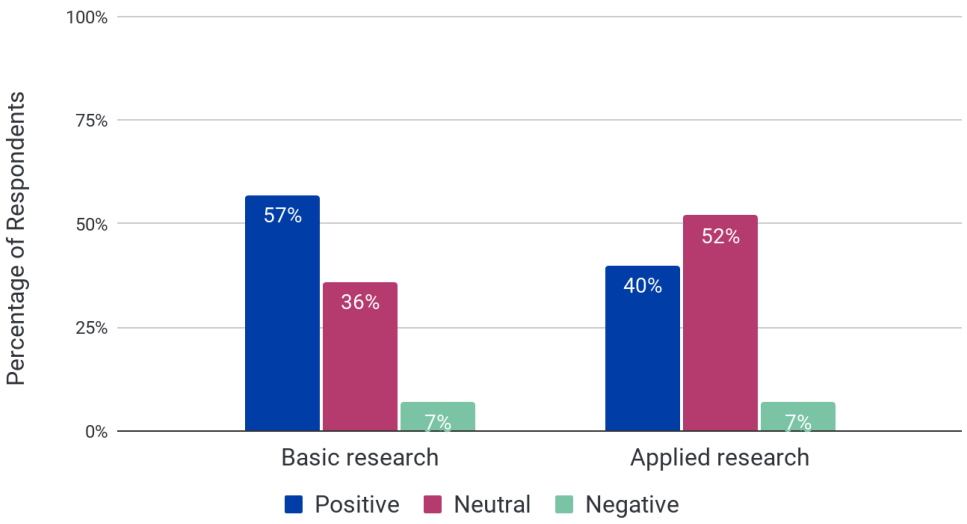

Question asked, "Do you generally feel negative, neutral, or positive about: accepting DOD grants for basic research / accepting DOD grants for applied research?" Respondents could select "generally negative," "neutral," or "generally positive." This question was only included in a pilot version of the survey. Respondents: 42. Proportions do not sum to 100 percent due to rounding. Source: CSET AI-DOD Survey 2020.

As displayed in Figure 3, some groups of respondents were more positive or negative in their views toward working on DOD-funded AI projects. Respondents with self-reported experience working at an employer engaged in DOD-funded research were more positive than respondents with no prior experience, especially if they worked directly on the DOD-funded project.[28] Respondents more familiar with DOD were also more positive, specifically those who indicated high familiarity with specific DOD AI projects, while respondents with less familiarity tended to be neutral. Lastly, respondents with high political trust were more positive about working on DOD grants and contracts, compared to respondents with little to no political trust. We explored several other subsample comparisons that resulted in no significant difference in positive or negative views toward working on DOD-funded AI projects, most notably respondent location on the East or West Coast and employer company size.[29]

Figure 3. AI Professionals who have Experience with DOD-Funded Work, High DOD Familiarity, and Political Trust Are More Positive About Working on DOD-Funded AI Projects

a. Experience with DOD-Funded Work

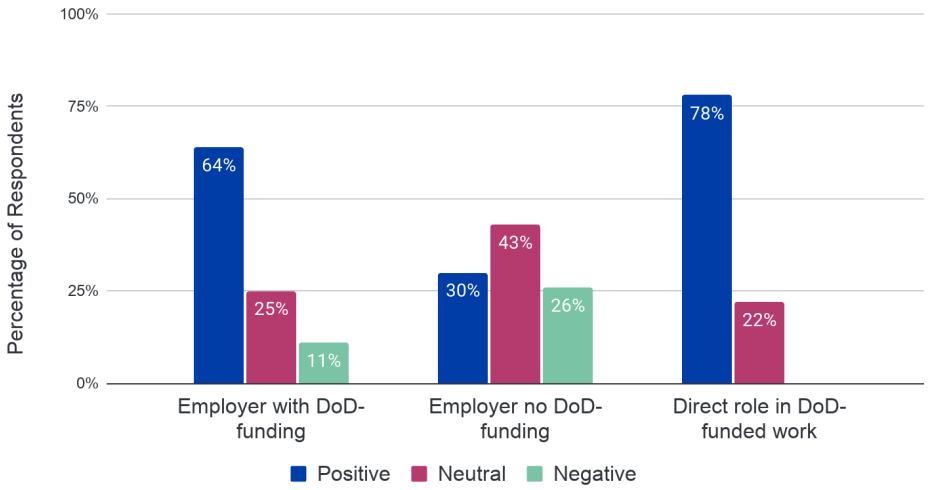

b. DOD Familiarity

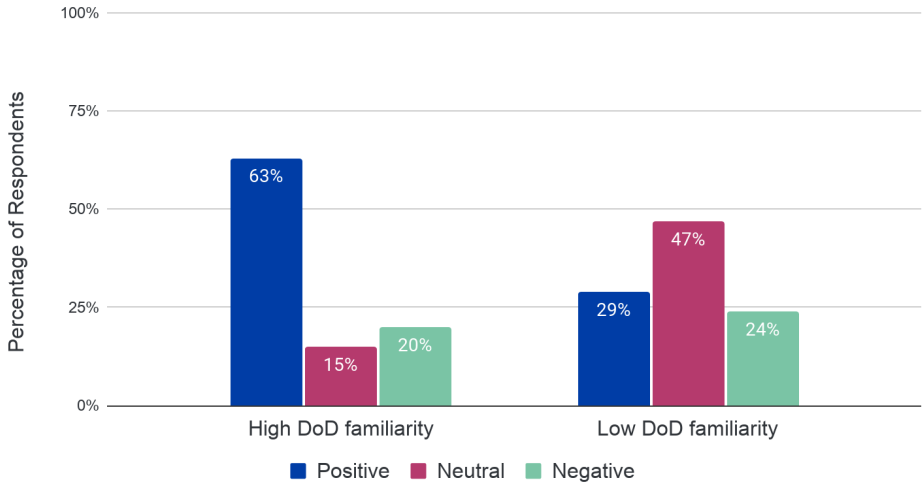

c. Political Trust

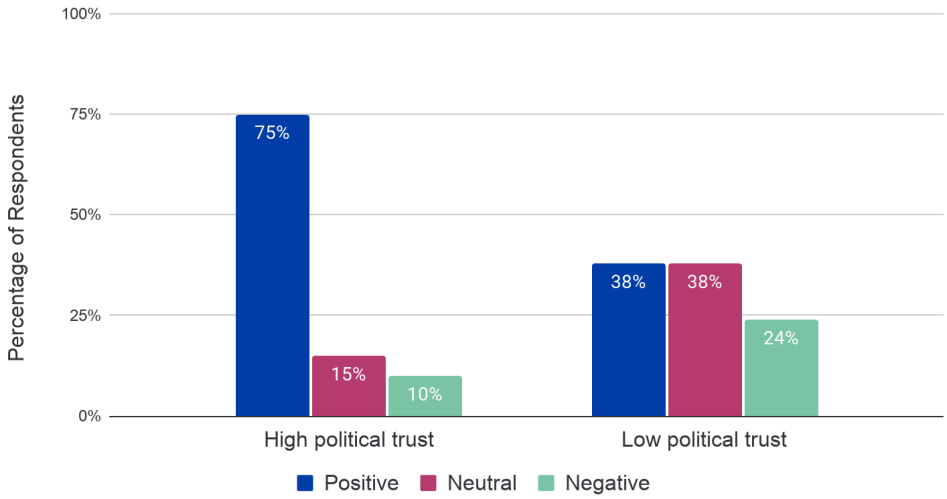

Percentage of respondents with positive, neutral, and negative views in each subsample. Respondents: experience with DOD-funded work - 59, direct role in DOD-funded work - 34, DOD familiarity - 74, political trust - 65. Source: CSET AI-DOD Survey 2020.

Respondents' general views about working on a DOD-funded AI project relate to other variables of interest. Respondents who are negative about working on DOD-funded projects rated funding source as a more important factor when deciding whether to work on a project than respondents who were positive.[30] Those who were positive also ranked defense-related U.S. government agencies more favorably than those who were negative toward working on DOD-funded research.[31]

Benefits and downsides to working on DOD-funded AI projects

As additional measures of AI professionals' attitudes toward DOD, we asked respondents what they consider to be the most compelling benefits and drawbacks to working on AI projects funded by DOD. First, we asked respondents to articulate the most convincing reasons for or against working on DOD-funded AI projects in their own words.[32] Then we asked them to select the benefits and drawbacks they find most compelling from a specified list of factors.

Benefits and reasons for working on DOD-funded projects

When asked to provide the most convincing reason to work on a DOD-funded AI project in their own words, about 65 percent of respondents said the ability to do good. Such responses included general statements about

doing good as well as specific statements about doing good for the country, the world, or science. Specific statements about doing good for the country were most common. For example, respondents said "helping protect my country and its citizens" or "preventing bad things from happening" were convincing reasons. Others said the good that comes from doing important research, working on "worthwhile project[s] with real international impact" and "basic research with wider positive impact" are convincing reasons.

Box 1. Reasons for Working on DOD-Funded AI Projects

Doing good -- "The DOD is in a unique position to bring about greater good than most organizations in the world today"

Doing good for country -- "Securing the safety of the United States against non-state and state threats."

Doing good for humanity -- "If the end results were proven to be for the advancement of all of humanity and the natural world we live in."

Advancing science -- "A lot of valuable things have come out of basic research that has been funded by the DOD. e.g., the internet, GPS and Kalman filters."

Research and funding opportunities -- "DOD is the only entity with deep enough pockets to fund big cutting edge research and development."

Forty percent of respondents said research and funding opportunities were convincing reasons for working on DOD-funded AI projects, especially the opportunity to work on interesting projects. For example, respondents said the most convincing reasons were the ability to work on "new challenges" and "basic research in an interest area" and access "technology that's not widely available (e.g. military drones)" and "unique datasets and computational resources." Others noted the funding attached to DOD grants and contracts, such as funding for "cutting-edge research" and "research which doesn't have an immediate commercial application." All of the reasons offered by respondents relate to doing good or research opportunities.

Figure 4 displays the results from the closed-ended question asking respondents what benefits are the most personally compelling from a specified list. Respondents were then given the option to rank selected benefits in order of importance. Interesting problems was the most frequently

selected benefit to working on DOD-funded AI projects and was also the top ranked benefit among a majority who selected it and provided a rank. A smaller majority selected access to unique resources as a compelling benefit. That these were the most frequently selected benefits reinforces findings from the open-ended responses, which also highlight research-related benefits. Other benefits such as civic duty and the ability to influence DOD were selected by just under half of respondents, again reinforcing the motivation of doing good for one's country we found in the open-ended responses.[33]

Figure 4. Interesting Problems Seen as Top Benefit to Working on DOD-Funded AI Projects

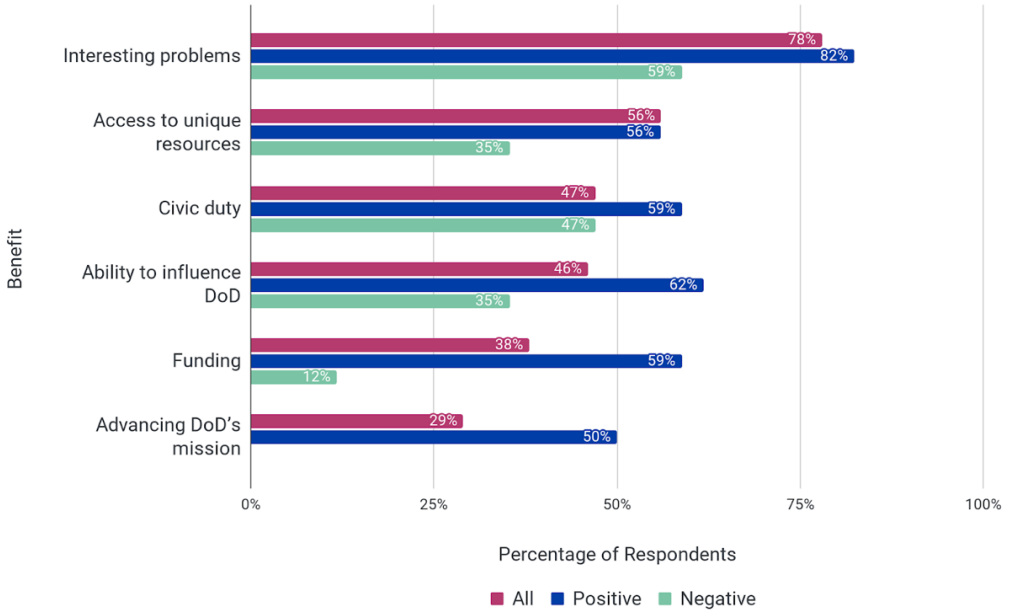

Question asked, "What do you see as the most personally compelling benefits of working on DOD AI grants or contracts?" Respondents could select all that apply. One respondent selected "other" and one respondent selected "none." Factors are listed on y-axis in order of selection by all respondents. Entries for "other" are listed in Appendix F. Respondents: all - 78, positive - 34, negative - 17. Source: CSET AI-DOD Survey 2020.

Perceptions of the most compelling benefits varied based on respondents' positive or negative view of working on DOD-funded AI projects. We find an especially large difference in the number of respondents with positive and negative feelings who selected advancing DOD's mission and funding as compelling benefits. No respondents who were negative about working on DOD-funded AI projects selected advancing DOD's mission as a benefit, compared to half of positive respondents who selected it. In other words, all respondents who selected advancing DOD's mission as a personally

compelling benefit are positive or neutral about working on DOD AI projects. The only benefit without a statistically significant difference in selection frequency between respondents with positive and negative views was civic duty.

Downsides and reasons for not working on DOD-funded projects

When asked to provide the most convincing reason not to work on a DOD-funded AI project in their own words, responses were more varied than reasons given for working on such a project. The most common downside, noted by 59 percent of respondents, was concern about doing harm. While this response was the most common, there was variation in the types of harm mentioned. Some offered general statements about not wanting to cause harm, such as concerns about "being involved in developing AI used for (physically) harming people" or working on capabilities that "may directly result in the death of other human beings." Others noted specific harm caused by war, autonomous weapons, mass surveillance, or a lack of ethics and safety protocols. Still others offered qualified harm-related reasons for not working on the project, such as if the project had a specific combat or offensive application.

Box 2. Reasons for Not Working on DOD-Funded AI Projects

Concerns about causing harm -- "Most DOD applications fundamentally have a high risk of harming human life whether directly (e.g. computer vision on drones) or indirectly (e.g. encouraging saber rattling instead of diplomacy). So, my bias is to assume the work would be used for purposes I don't agree with and would need to be convinced that it is in one of the nooks and crannies at the DOD where that would not be the case."

Concerns about taking human life -- "While effective defense and monitoring capability are important, I have a strong impression that a lot of DOD work at this point is basically pointless murder and I would not participate in that under any circumstances."

Concerns about safety -- "DOD stakeholders who don't understand the potential and limits of AI"

Concerns about use -- "[I] don't want to develop weapons or surveillance technology."

Box 2 continued.

Disagreement with DOD's mission -- "The end purpose of DOD is not one of bringing humanity together but rather based on the premise that our nations are separate, unequal, and do not have an underlying unity as human beings. Though some of the research and tech can be used for 'social good' this is not the ultimate goal of DOD."

Research constraints -- "If the results turn out to be not what is expected or required, the whole work will be wasted...A DOD contract must have very specific outcomes."

The next most common reasons given for not working on a DOD-funded AI project related to research constraints, such as publications restrictions, monetary issues, or bureaucratic requirements. For example, respondents mentioned "[the] inability to publish results" and "not being able to have full control over the direction of the project" as reasons not to work on DOD-funded AI projects. Seven respondents explicitly said there are no convincing reasons for not working on a DOD-funded AI project.[34]

Figure 5 displays the results for the closed-ended question asking respondents what downsides to working on a DOD AI project are most personally compelling from a specified list. Again, respondents then had the option to rank their selected downsides. Discomfort with how DOD might use the work was the most frequently selected downside and was ranked as the most important downside for a majority of respondents who selected it and provided a ranking. As we found with open-ended responses, this downside stands out as most common among a list of varied other downsides. Other frequently selected downsides relate to research restrictions, including paperwork and reporting requirements, security clearance and classification requirements, under-resourced projects, and publishing restrictions. Less commonly selected downsides include disagreement with DOD's mission, negative reactions from peers, and low prestige work.

Figure 5. Discomfort with Use is Major Downside to Working on DOD-Funded AI Projects

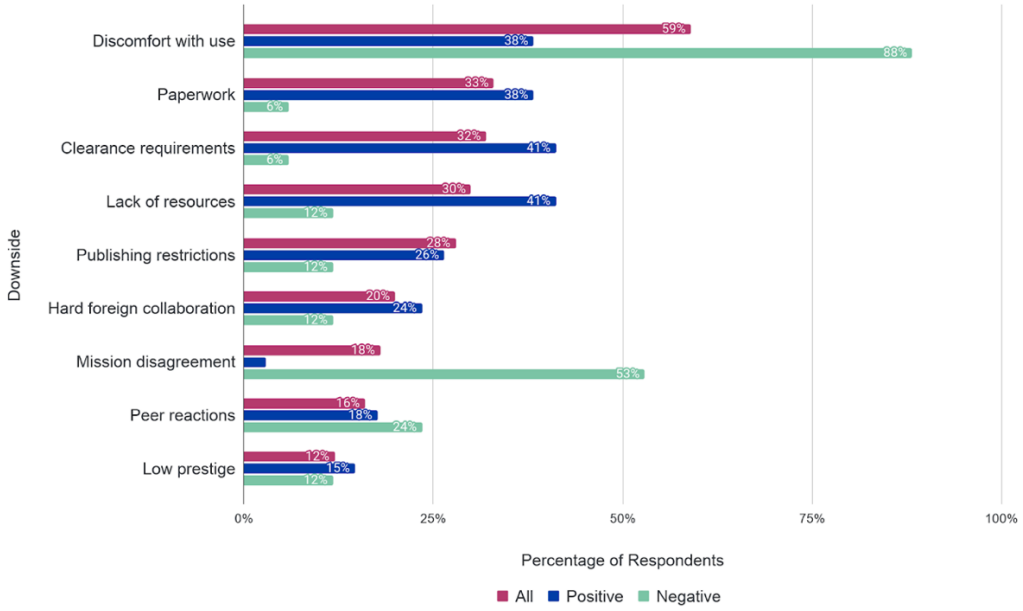

Question asked, "What do you see as the most personally compelling downsides to working on DOD AI grants or contracts?" Respondents could select all that apply. Differences in proportions for positive and negative respondents is statistically significant at 95 percent confidence level for discomfort with use, paperwork, clearance requirements, under-resourced project, and disagreement with DOD's mission. Factors are listed on y-axis in order of selection by all respondents. Entries for "other" are listed in Appendix F. Respondents: all - 76, positive - 34, negative - 17. Source: CSET AI-DOD Survey 2020.

Again, we find differences in the downsides selected based on respondents' positive or negative view of working on DOD-funded AI projects. Respondents with a negative view overwhelmingly selected discomfort with how DOD will use the work and much more commonly selected disagreement with DOD's mission as personally compelling downsides. In comparison, a minority of respondents with a positive view selected discomfort with how DOD will use the work as a compelling downside and only one respondent with a positive view selected disagreement with DOD's mission as a downside. Meanwhile, respondents who feel positively about working on DOD-funded AI projects more often selected under-resourced projects and clearance restrictions as compelling downsides, while these factors were selected by only six percent of respondents with negative views. This suggests that downsides that are considered compelling are related to respondents' general views about working with DOD. Those with negative general views find their discomfort and disagreement with DOD to be a more compelling

reason not to work on DOD-funded AI projects. Those with positive general views are more varied in the downsides they find compelling, but are more convinced by research-related downsides than by discomfort or disagreement.

Factors that would increase willingness to work on DOD-funded projects

In addition to asking respondents the perceived benefits and downsides of working on DOD-funded AI projects, we asked what factors would increase their willingness to work on DOD-funded AI projects. Respondents could choose from a specific list of nine factors plus write in any additional factors. Results are displayed in Figure 6. The most frequently selected factor that would increase willingness was the ability to provide a global benefit (e.g., humanitarian relief). A majority of respondents also said willingness would increase if funding was for basic research only. Just over 40 percent said projects with defensive (rather than offensive) applications or a stronger commitment by DOD to safety and security or privacy and bias risks would increase their willingness. Fewer respondents selected increased threat to the United States from foreign adversaries, reduced publication restrictions, or a project done in collaboration with U.S. allies as factors that would increase willingness.

Figure 6. Providing Global Benefit Increases Willingness to Work on DOD-Funded AI Projects

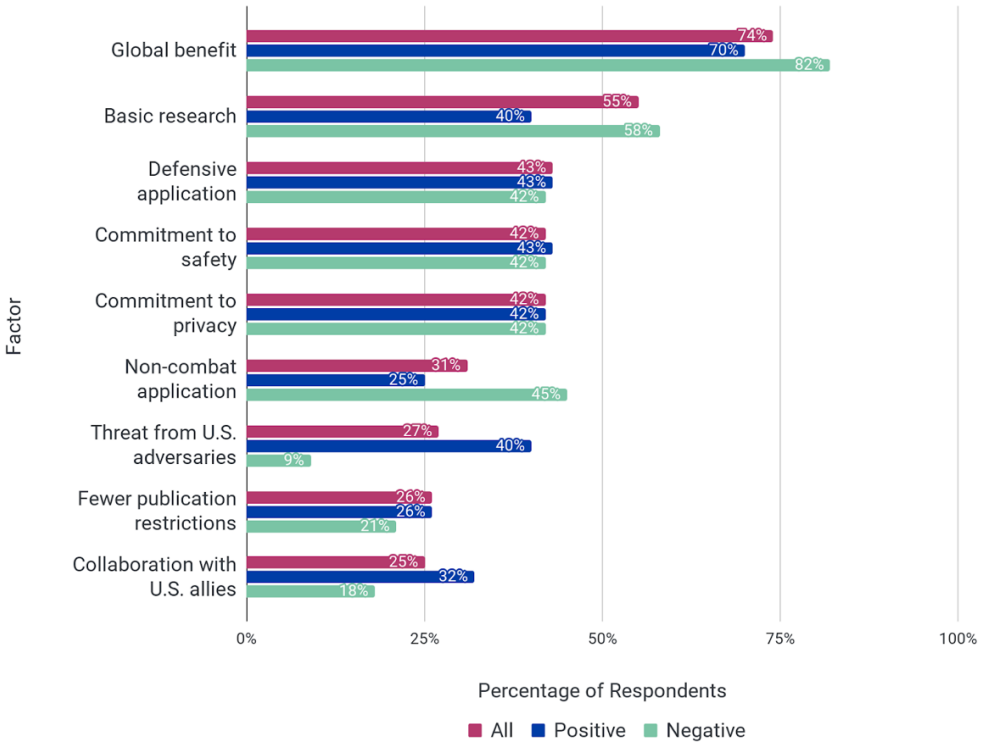

Question asked, "Which of the following factors would increase your willingness to work on a DOD AI grant or contract?" Respondents could select all that apply. Factors are listed on y-axis in order of selection by all respondents. Entries for "other" are listed in Appendix F. Respondents: all - 133, positive - 53, negative - 33. Source: CSET AI-DOD Survey 2020.

After selecting the factors that would increase their willingness, respondents were asked to indicate the amount that those factors would increase willingness. Among the majority who selected global benefit, most said that factor would increase willingness a lot. Similarly, a majority of respondents who selected funding for basic research and stronger commitments to privacy and bias risks said they would increase willingness a lot.

Comparing the factors that would increase willingness for respondents with positive and negative views about working on DOD-funded AI projects, we find fewer differences than when comparing benefits and downsides. The ability to provide a global benefit was the most selected factor regardless of baseline positive or negative views about working on DOD-funded AI projects. This finding suggests that across the board, AI professionals are more willing to work on DOD-funded AI projects that provide for global humanitarian benefit.

Two factors have significant differences across respondents with positive and negative views of working on DOD-funded AI projects. First, 45 percent of respondents with negative views selected administrative or non-combat applications as a factor that would increase willingness, making it the third most selected factor among that group. A significantly smaller 25 percent of respondents with positive views selected it; it was the least selected factor among those with positive views. Second, 40 percent of respondents with positive views selected increased threats to the U.S. from foreign adversaries as a factor that would increase willingness, but only nine percent of respondents with negative views selected it. These differences suggest respondents are differentially motivated by these specific factors: willingness among those who already feel positive about working with DOD is increased by foreign threats, while willingness among those who already feel negative is increased by working on projects without combat applications.

We also asked respondents whether DOD's adoption of AI ethical principles in 2020 changed their willingness to work on DOD-funded AI projects. Results are displayed in Figure 7. Respondents were first asked if they were familiar with the newly adopted principles and were provided a link to review the principles. Most, 67 percent, were not at all familiar with the principles. When asked if DOD's adoption of a set of ethical principles changed their willingness, over half said it did not change their willingness. Of those who said it did change their willingness, almost all said it increased their willingness to work on DOD-funded AI projects.

Figure 7. DOD's AI Ethics Principles Have Either No Impact or Positive Impact on Willingness to Work on DOD-Funded AI Projects

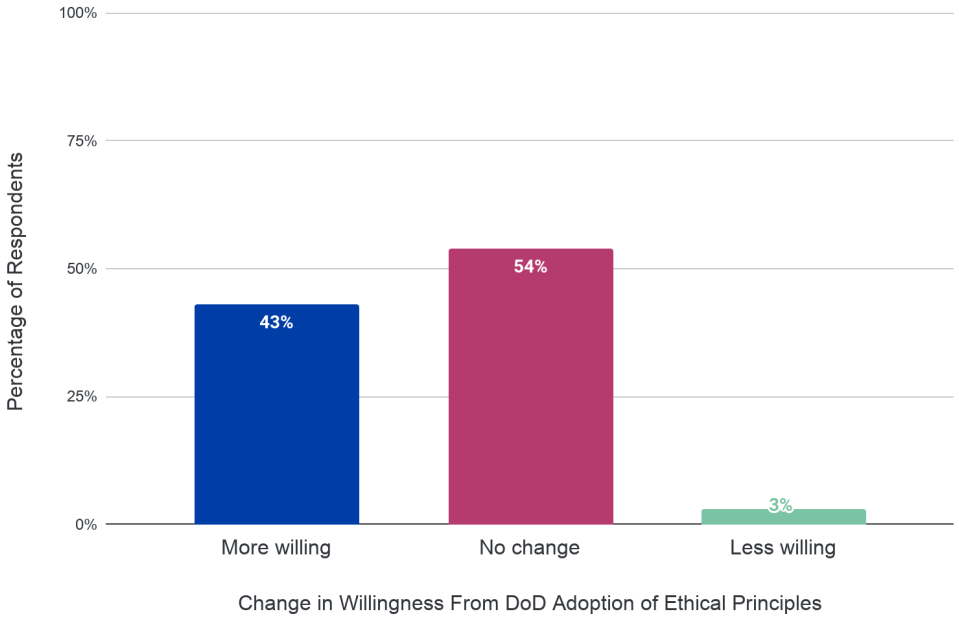

Question asked "How does your willingness to work on DOD AI grants or contracts change based on the adoption of these principles?" and respondents could select "more willing," "no change," or "less willing." Respondents: 72. Source: CSET AI-DOD Survey 2020.

Impact of project application and scope on willingness and actions

The survey included an experiment to test the impact of project application and scope on respondents' attitudes toward DOD-funded AI projects. The experiment described a hypothetical DOD AI contract and asked respondents about their willingness to work on the contract and what actions they would take in response to their employer bidding on the contract. The contract varied along two dimensions, end-use application and project scope, resulting in four treatment conditions. A fifth condition presented a DOD contract to implement AI in back-office department administration.

Box 3. DOD Contract Descriptions for Experimental Conditions

U.S. Humanitarian -- As part of a larger initiative to assist U.S. disaster relief efforts, a DOD contract provides funding for a project to apply machine learning capabilities to aid in the prediction of extreme weather events. Your company has relevant expertise and considers putting in a bid for the contract.

Global Humanitarian -- As part of a collaboration with U.S. allies to assist global disaster relief efforts, a DOD contract provides funding for a project to apply machine learning capabilities to aid in the prediction of extreme weather events. Your company has relevant expertise and considers putting in a bid for the contract.

U.S. Battlefield -- As part of a larger initiative to assist U.S. combat efforts, a DOD contract provides funding for a project to apply machine learning capabilities to enhance soldier effectiveness in the battlefield through the use of augmented reality headsets. Your company has relevant expertise and considers putting in a bid for the contract.

Global Battlefield -- As part of a larger initiative with U.S. allies to enhance global security, a DOD contract provides funding for a project to apply machine learning capabilities to enhance soldier effectiveness in the battlefield through the use of augmented reality headsets. Your company has relevant expertise and considers putting in a bid for the contract.

Department Back-office -- A DOD contract provides funding for a project to apply machine learning capabilities to enhance the DOD's back-office operations. Your company has relevant expertise and considers putting in a bid for the contract.

Willingness to work on DOD AI contract

We found that willingness did significantly vary across the experimental conditions: respondents were more willing to work on humanitarian projects, as opposed to battlefield projects. We did not find that project scope significantly altered willingness to work on the project.

Across conditions, 64 percent of respondents were willing to work on the described contract, including 41 percent who were very willing to work on the project. This suggests that unwillingness is not the default position of AI professionals when it comes to working on a range of DOD contracts.

Figure 8 displays the percentage of respondents who indicated a willingness to work on the described contract in each condition. Respondents were most willing to work on the U.S.-focused humanitarian project, with 86 percent of respondents in that condition willing to work on the project. A comparable 81 percent were willing to work on the global humanitarian project in collaboration with international allies. Meanwhile, only 42 percent of respondents in the U.S.-focused battlefield condition were willing to work on the project.

Figure 8. Willingness to Work on DOD-Funded AI Projects is Higher for Humanitarian Projects

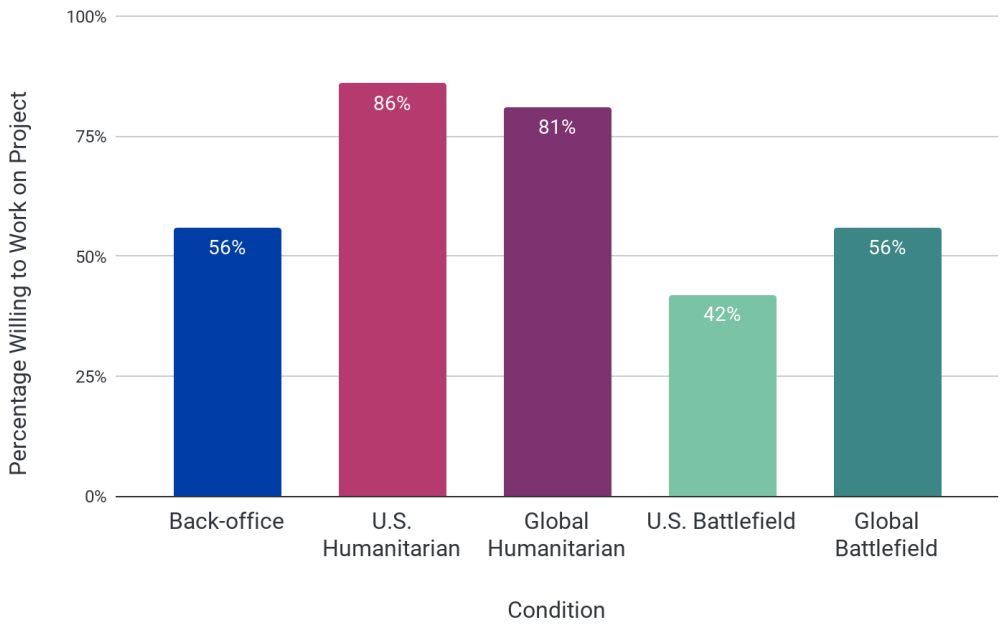

Percentage of respondents who selected "very willing" or "somewhat willing" in response to "please indicate your degree of willingness to work on this project." Respondents: back-office - 25, U.S. humanitarian - 21, U.S. battlefield - 19, global humanitarian - 26, global battlefield - 27. Source: CSET AI-DOD Survey 2020.

Figure 9 shows the differences in mean willingness in different experimental conditions. The most significant difference in willingness was comparing humanitarian to battlefield project applications. Average willingness was significantly higher in humanitarian conditions, compared to the battlefield conditions.[35] Willingness was also higher in the humanitarian conditions than in the department back-office condition.[36] We found no significant difference in mean willingness between the U.S.-focused and global conditions.

Figure 9. Mean Willingness Is Significantly Higher for Humanitarian Projects Compared to Battlefield Projects

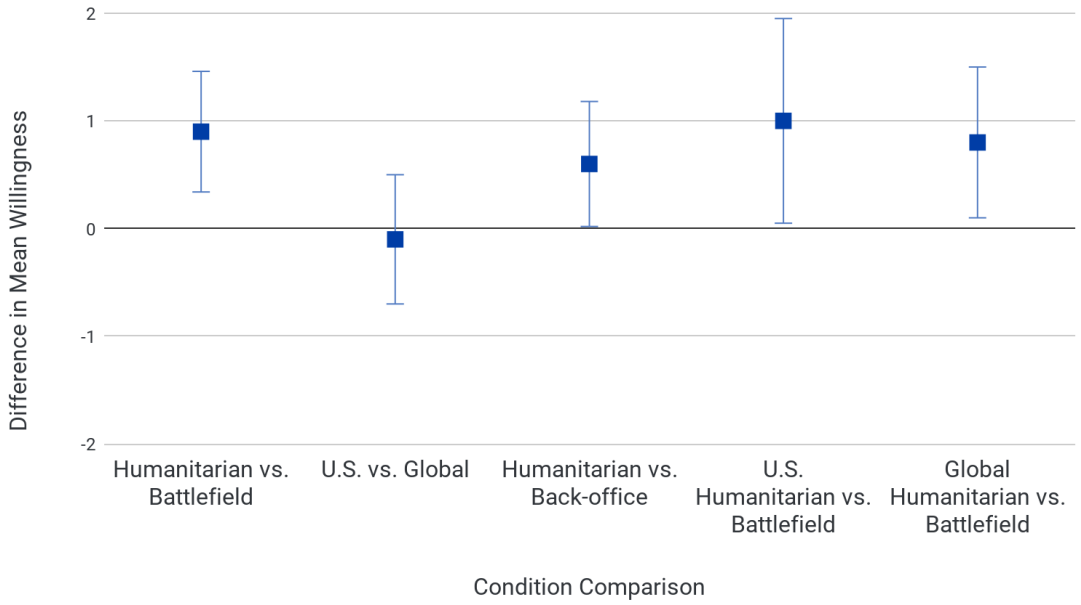

The difference in mean willingness for select project condition comparisons. Error bars provide a 95 percent confidence interval around the estimated difference. Where the confidence intervals do not cross zero, we can reject the null hypothesis of no difference between the two conditions. Source: CSET AI-DOD Survey 2020.

That willingness to work on DOD-funded AI projects was not influenced by whether the project had a U.S. or global application is surprising given respondents' frequent selection of global benefit as a factor that would increase willingness to work on a DOD-funded AI project. These different findings are likely the product of survey design. The relevant response option for the factors that would increase willingness was worded as the "ability to provide global benefit (e.g., humanitarian relief)," which could be read as combining global scope and humanitarian application. This might suggest that respondents consider providing global humanitarian relief as a compelling reason to work on DOD-funded AI projects, but that it is the distinction between humanitarian and non-humanitarian applications, as opposed to global versus U.S.-specific scope, that is increasing willingness.

Actions considered in response to employer bidding on DOD contract

We also asked respondents what actions they would take, if any, if their employer were to bid on the described DOD contract.[37] Among all respondents, the most common actions considered were expressing interest in

being on the project (36 percent), taking no action (24 percent), and avoiding working on the project (20 percent). We found a significant difference in the actions respondents considered in the humanitarian and battlefield conditions, but again found no difference in the actions considered based on U.S. as opposed to global scope.

Figure 10 displays the actions considered by respondents in each condition. The most common action considered for humanitarian projects was expressing interest in being on the project, while the most common action considered in the battlefield conditions was to avoid working on the project.[38] The frequency of respondents considering expressing interest in humanitarian conditions was higher than the frequency of respondents avoiding the project in battlefield conditions, suggesting respondents were more willing to proactively engage with projects they support rather than actively condemn projects they do not support. Another difference to note is that there is more variation in the actions considered in battlefield and department back-office conditions, while expressing interest in working on the project and taking no action are most prevalent in the humanitarian conditions.

Figure 10. AI Professionals Consider Expressing Interest in Humanitarian Projects and Avoiding Battlefield Projects.

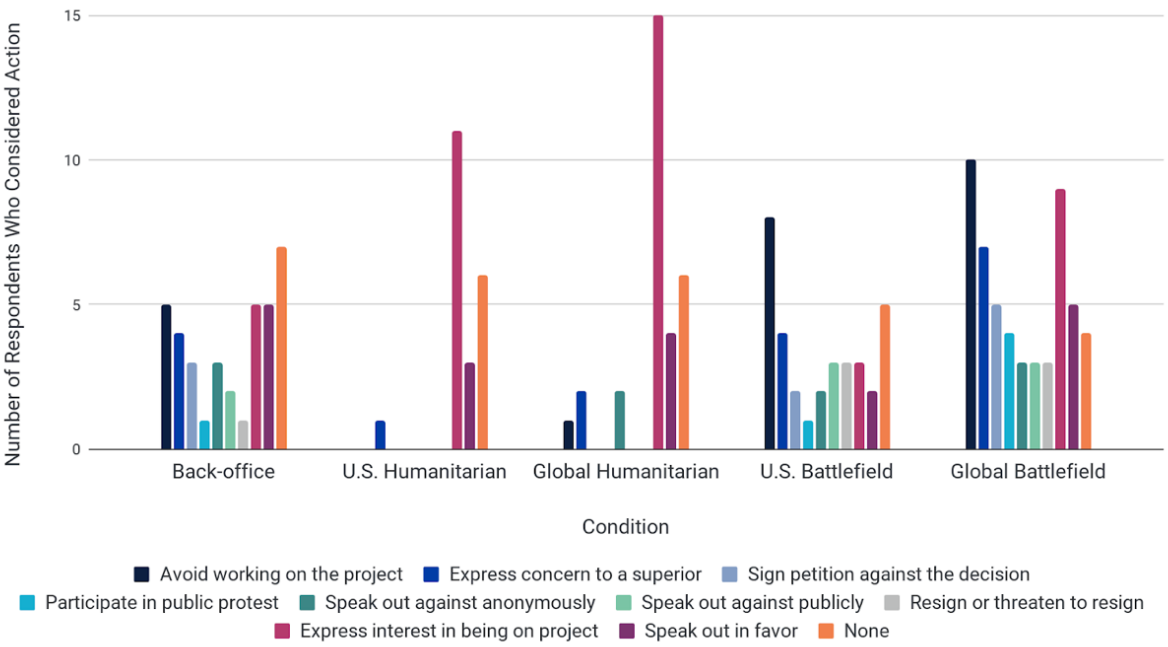

Count of actions considered by respondents in each experimental condition. Respondents: back-office - 25, U.S. humanitarian. - 21, U.S. battlefield. - 19, global humanitarian - 26, global battlefield - 27. Source: CSET AI-DOD Survey 2020.

Other actions included in the survey were less frequently considered. Some considered expressing concern to superiors about the project, especially in the battlefield conditions, or speaking out in favor of the decision to bid on the contract. The actions considered by the fewest respondents were participating in public protest, resigning or threatening to resign, and speaking out publicly against the decision to bid on the contract.

Figure 11 compares the proportion of actions considered that are in support of the employer's decision to bid on the DOD contract to the actions that are against the employer's decision. Again, we find most actions considered in the humanitarian conditions are in support of the bid, either expressing interest in working on the project or speaking out in favor of the bid, especially for the U.S. humanitarian project. Meanwhile the majority of actions considered in the battlefield and even department operations conditions are in protest of the decision. Note that the survey included more actions against the employer's decision which range in their degree of protest, from resigning to speaking out anonymously to avoiding the project.

Avoiding the project and expressing concern to a superior were the most frequently selected protest actions.

Figure 11. AI Professionals Consider Supportive Actions for Humanitarian Projects

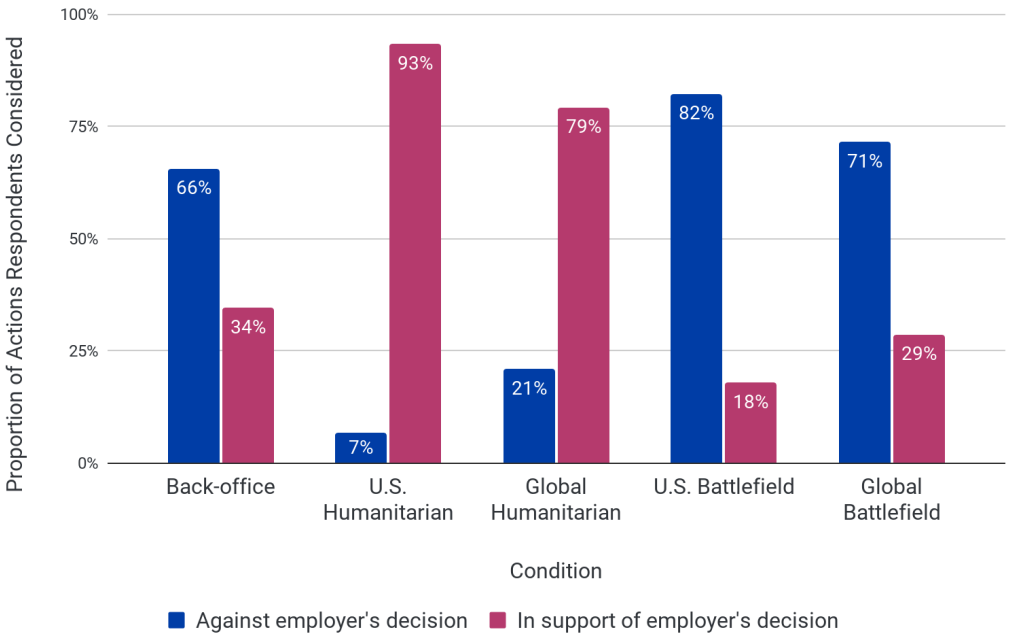

Percentage of actions considered that were in support of or against employer decision to bid on the contract in each experimental condition. Actions included as against employer's decisions are avoiding working on the project, expressing concern to a superior, signing a petition against the bid, participating in public protest, speaking out against the bid publicly or anonymously, and resigning. Actions included as in support of employer decision are expressing interest in working on the project and speaking out in favor of the bid. Respondents: back-office - 25, U.S. humanitarian. - 21, U.S. battlefield. - 19, global humanitarian - 26, global battlefield - 27. Source: CSET AI-DOD Survey 2020.

Experience and familiarity with DOD

We asked respondents whether they had previous professional experience working on DOD grants or contracts. Just over half of respondents work or have worked for an organization with DOD-funded contracts.[39] In more than half of those cases, 56 percent, it was in a past position. Of those who reported working for an organization engaged in DOD-funded work at any point, half said they were directly involved in the funded work. Comparing responses to data on DOD contracts awarded since 2015, we find similar results. Respondents currently work at 62 companies and 47 percent of those companies have had a DOD contract since 2015.[40] That translates to 63

percent of respondents who currently work at a company that has had a DOD contract.

We asked several additional questions to gauge respondents' familiarity with DOD. As seen in Figure 12, respondents had limited familiarity with specific information about DOD. A large majority were not at all familiar with DOD's budget categories, 67 percent were not at all familiar with DOD's AI ethical principles, and 59 percent were not at all familiar with specific DOD-funded AI projects. Respondents were more familiar with the defense-related, but not DOD-specific, debate around the development of lethal autonomous weapons.

Figure 12. Respondents Are Not Very Familiar with DOD Projects and Terminology

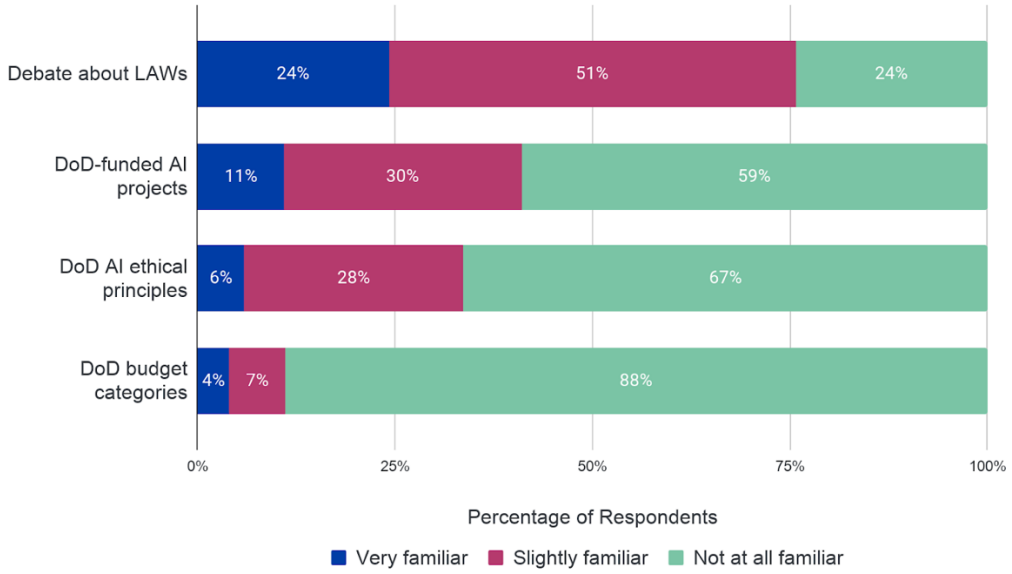

Questions asked "In February 2020, the DOD adopted a set of AI ethical principles. How familiar are you with the newly adopted AI ethical principles?" and "Indicate your familiarity with the following: Specific DOD-funded AI projects (e.g. Army's ATLAS program), the DOD budget categories (6.1-6.7), international debate on the development of lethal autonomous weapons." Respondents: 70. Source: CSET AI-DOD Survey 2020.

Professional considerations

We also asked respondents several questions about how they approach their professional obligations and assess ethical considerations around technological development. First, we asked respondents whether concerns about working with DOD have ever affected their decisions about where to work, what research or projects to work on, what colleagues to work with, or

whether to persuade others not to work on specific projects. As displayed in Figure 13, the most frequent impact of DOD concerns was on past decisions about where to work or what projects to work on.

Figure 13. DOD Concerns Impacted Some AI Professionals' Decisions About Where to Work

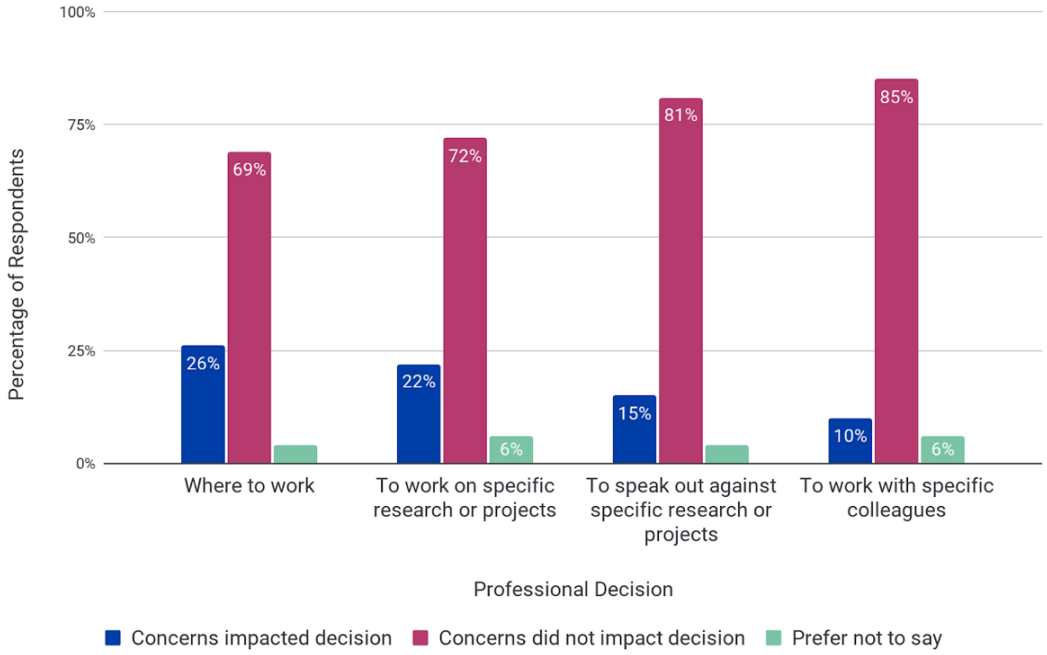

Question asked "Have concerns about working with DOD ever affected your decision(s) about: Where to work? Whether to work on specific research or projects? Whether to work with specific colleagues? Whether to speak out or persuade others not to work on specific research or projects?" Respondents: 72. Source: CSET AI-DOD Survey 2020.

Next we asked respondents their level of agreement with various statements about the consequences of their work. As displayed in Figure 14, respondents agreed with the idea that they, and their colleagues, regularly think about the implications of their work. There was almost unanimous agreement that respondents regularly think about the potential consequences of the technology they work on, including 49 percent of respondents who strongly agree. There was also high agreement that social impact is a major decision factor in whether to work on a project and that colleagues tend to agree on issues about the implications of their work.[41]

Figure 14. AI Professionals Think About the Consequences of Their Work and Say It Factors into Projects They Work On

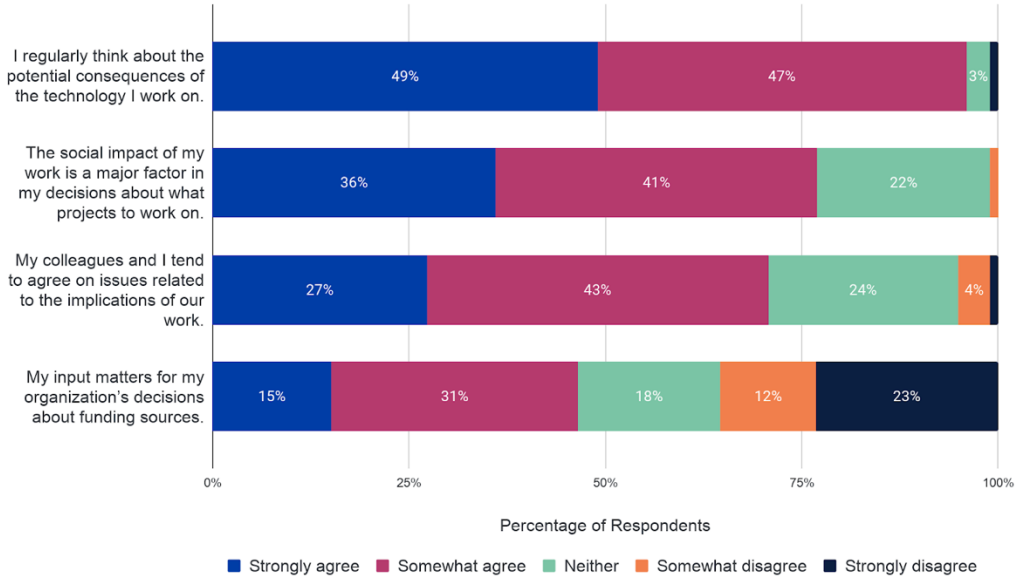

Question asked respondents to indicate agreement with the listed statements. Respondents: 74. Source: CSET AI-DOD Survey 2020.

Finally, we asked respondents how they think about the sources of funding for projects they work on. Most respondents are not usually aware of the funding source for projects they work on. Only 17 percent report always knowing project funding source(s) and a combined 31 percent report knowing project funding source(s) more than half of the time. For respondents who are aware of funding sources for projects they work on, 56 percent consider funding source(s) to be moderately, very, or extremely important in their decision whether to work on the project. Of those who consider project funding source important, 54 percent said the mission of the funder had the greatest impact on their preferences about that funder. In addition to being unaware of funding sources, respondents are split on whether their input matters for company decisions about funding decisions: 46 percent say it does while 35 percent disagree.[42]

AI development

We were also interested in AI professionals' views on AI development. To explore these views, we asked respondents how much they trust various actors to develop AI in the best interest of the public. As displayed in Figure 15, the most trusted actors were academia and the respondent's employer, both trusted a fair amount or a great deal by more than 75 percent of

respondents. The least trusted actor was the Chinese government, which most respondents do not trust at all. Intergovernmental organizations (IGOs) and the European Union were ranked as more trustworthy than U.S. tech companies, the U.S. government, and the U.S. military – each of which was ranked as trustworthy by a minority of respondents. A 2019 survey of AI/ML researchers found similar levels of trust in these actors to develop and manage AI in the best interests of the public.[43]

Figure 15. Employer, Academia, and European Union Considered Most Trustworthy to Develop AI

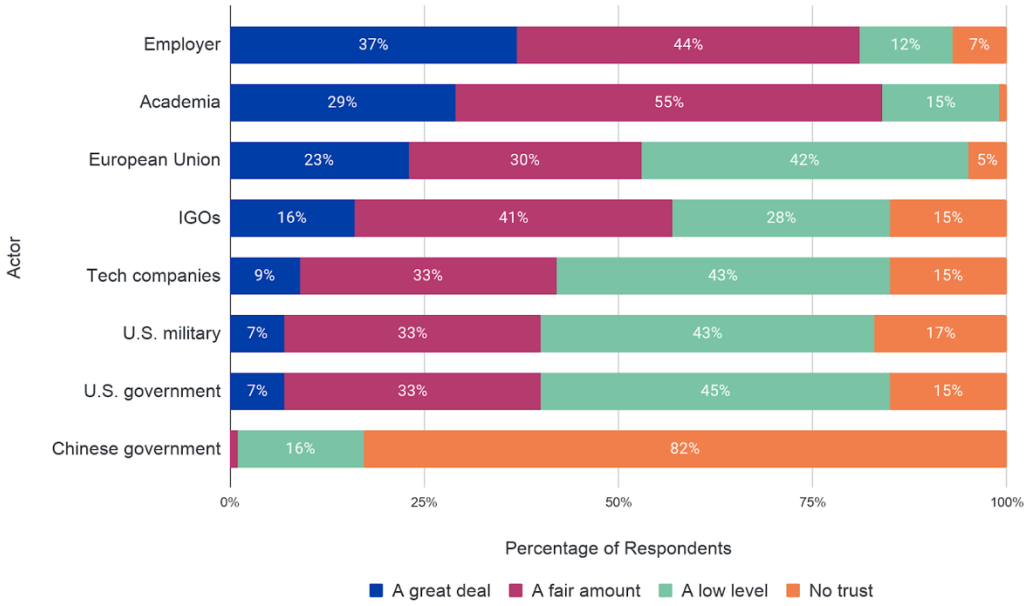

Question asked "Indicate your level of trust in each of the following organizations to develop AI in the best interests of the public." Respondents provided their level of trust for each listed organization. Respondents: 75. Source: CSET AI-DOD Survey 2020.

Comparing our results to a 2018 survey of Americans, we find some differences in levels of trust in actors to develop AI, displayed in Figure 16.[44] First, the sample of the American public is relatively confident in the U.S. military's ability to develop and manage AI in the interests of the public, second only to confidence in university researchers by a small margin. In contrast, our sample of AI professionals rank several actors as more trustworthy than the U.S. military in terms of developing AI in the interests of the public. Yet, our sample of AI professionals considers the U.S. government slightly more trustworthy. Both samples considered academia to be a trustworthy actor in terms of developing AI, but our sample of AI professionals has a higher degree of trust in academia.[45]

Figure 16. AI Professionals are More Trusting of Academia and Intergovernmental Organizations to Develop AI and American Public is More Trusting of U.S. Military to Develop AI

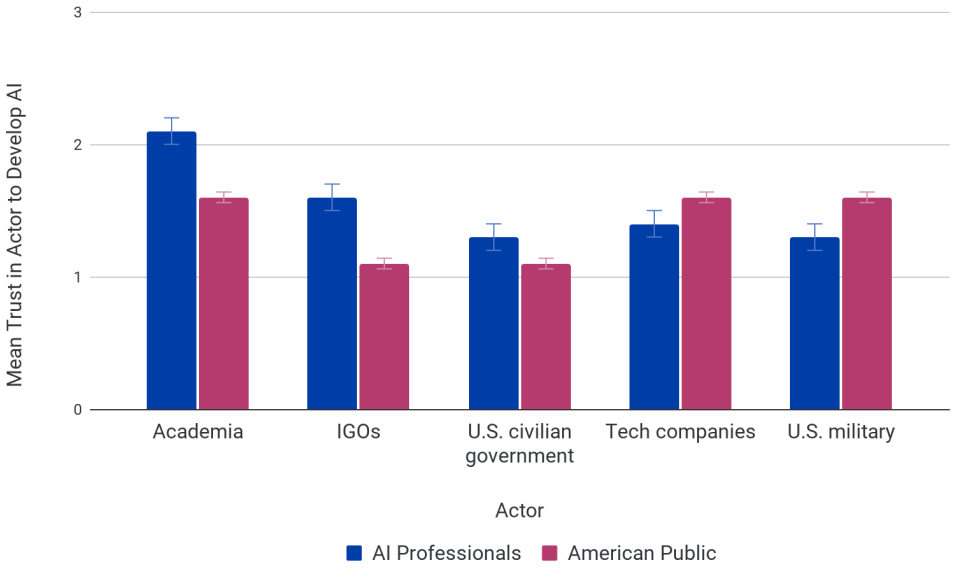

Mean trust in various actors to develop AI in the best interests of the public where 0 is "no trust/confidence," 1 is "a low level of trust/not too much confidence," 2 is "a fair amount of trust/confidence" and 3 is "a great deal of trust/confidence." For AI professionals, the question asked "Indicate your level of trust in each of the following organizations to develop AI in the best interests of the public." Respondents provided their level of trust for each of eight actors. Respondents: 75. Source: CSET AI-DOD Survey 2020. For the American public, the question asked "How much confidence, if any, do you have in each of the following to develop AI in the best interests of the public?" Respondents were shown five randomly selected actors from a list of 14 actors. Mean trust was calculated using raw frequencies and excludes don't know responses. Source: Center for the Governance of AI 2019.

Lastly, we asked respondents what actors should be responsible for mitigating negative consequences of new AI technologies and to rank those actors in terms of responsibility. Nearly all respondents said national governments should be responsible for managing consequences of new AI tech and ranked them as the most responsible actor. Ninety-one percent said companies should be responsible, but respondents tended to rank companies as less responsible than national governments. A majority of respondents considered individual researchers, managers, or IGOs responsible, but when these actors were considered responsible, IGOs were ranked as more responsible than managers or individual researchers, who were ranked least responsible. The actors most frequently considered not responsible were managers and individual researchers.

Figure 17. Respondents Consider National Governments and Companies Responsible for Mitigating Negative Consequences of AI Technology

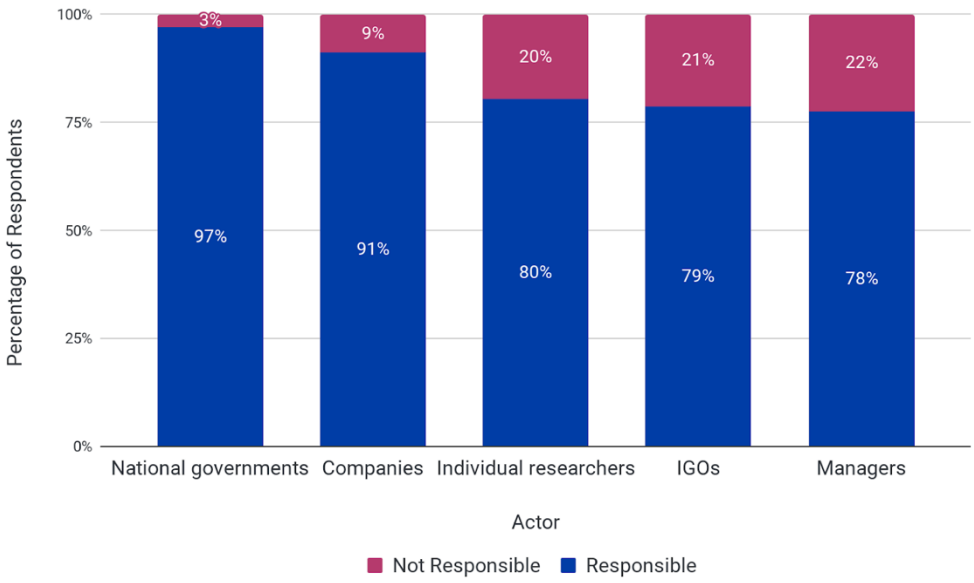

Question asked "Indicate which of the following actor(s) should be responsible for mitigating negative consequences of new AI technologies and rank them in terms of their level of responsibility." Respondents: 62. Source: CSET AI-DOD Survey 2020.

Political considerations

Our final questions asked respondents about their political activities and attitudes toward the U.S. government. First, we asked respondents to rate five U.S. federal government departments and agencies in terms of favorability. As displayed in Figure 18, respondents felt most favorable toward the National Science Foundation (NSF) and the National Institute of Standards and Technology (NIST). A majority of respondents rated both favorably. A majority also rated the Defense Advanced Research Projects Agency (DARPA) favorably.

Figure 18. Respondents View NSF, NIST, and DARPA Favorably and DHS Unfavorably

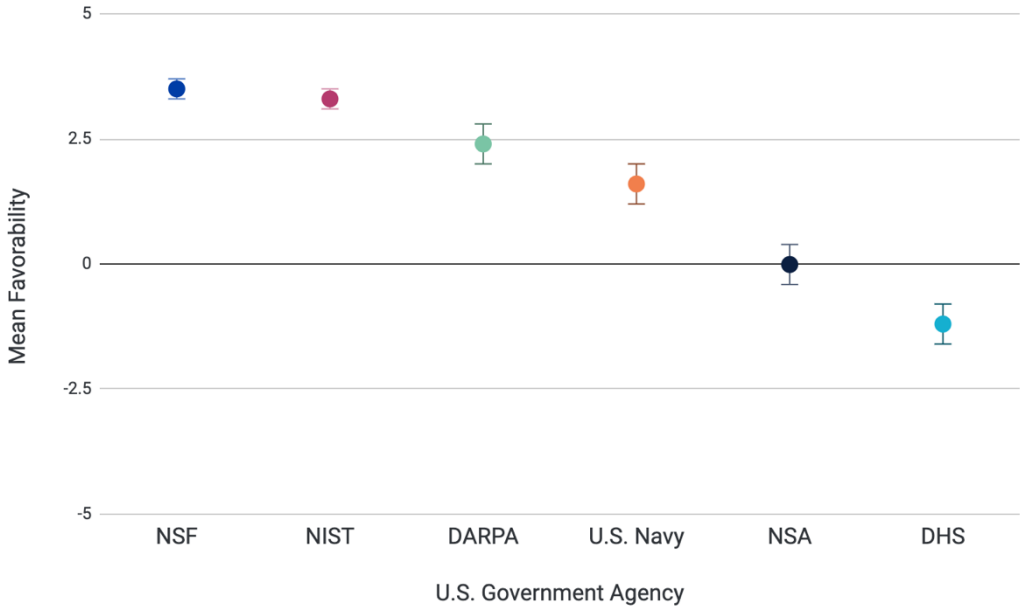

Question asked "Please indicate your feelings toward the following U.S. government agencies using the scale provided." A mean favorability between 0 and 5 indicates positive feelings and a mean favorability between 0 and -5 indicates negative feelings. A mean favorability of 0 indicates neutral feelings toward the agency. All respondents rated NSA, U.S. Navy, DARPA, and DHS and were randomly assigned to rate either NSF or NIST. Respondents: 67. Source: CSET AI-DOD Survey 2020.

Respondents were least favorable toward the Department of Homeland Security (DHS). There was greater variability in ratings of the Department of the Navy and the National Security Agency but the U.S. Navy was mostly rated favorably while NSA ratings skewed slightly toward unfavorable.

Respondents were politically interested and involved. Seventy-one percent were very interested in what's going on in government and politics; 35 percent have donated to a political campaign in the last year. Respondents were also relatively trusting of the U.S. government, with 31 percent saying you can trust the U.S. government to do what is right most or all of the time. Based on these estimates, our sample is more politically interested, involved, and trusting than average Americans.[46] This is not entirely surprising given the known link between socioeconomic status and political participation and interest. This population has higher levels of education and income than the general public. Yet, finding they are also more trusting is interesting because the link between socioeconomic status and political trust is not as clear.

Conclusion

In August 2020, Palantir CEO Alexander Karp re-affirmed the company's support of DOD, drawing a sharp contrast between the company's values and those presumed held by the broader tech industry. "[Palantir] seem[s] to share fewer and fewer of the technology sector's values and commitments," he lamented before pledging Palantir's unwavering support of U.S. defense: "we have chosen sides."[47] His comments echo the popular narrative of a tech industry in unified opposition to U.S. defense, but how accurate is that narrative? When it comes to the use of computing technologies such as AI for national defense, has the tech industry chosen a side?

Our findings suggest that is not the case: surveyed AI professionals hold a range of views toward DOD. Many AI professionals, meaning professionals with AI-skills employed at AI companies in AI hubs, hold no strong feelings about working on DOD-funded AI projects, and only a minority hold extreme views, positive or negative. We also found that views vary depending on the nature of the project. AI professionals are more open to working on DOD-funded AI projects if they have humanitarian applications or involve basic research. AI professionals' top concerns about working on DOD-funded AI projects relate to how the technology might be used and they articulate concerns about the potential harm caused by AI technologies.

Despite these similarities, motivations for working on DOD-funded AI projects vary depending on general views toward DOD and the U.S. government. AI professionals tend to be more positive about working on DOD-funded AI projects if they are more familiar with DOD, have experience working on DOD-funded projects, or have higher trust in the U.S. government. Conversely, those less familiar with DOD or working on DOD-funded projects tend to be more negative about working on DOD-funded AI projects. AI professionals who feel negatively about working on DOD-funded AI projects are more interested in projects that support basic research and non-combat applications. Relative to those who feel negatively, AI professionals who feel positively about working on DOD-funded AI projects find the prospect of influencing DOD and gaining access to unique resources and funding as more compelling benefits to working on such projects.

This survey is a first step toward an evaluation of the current relationship between the tech industry and DOD. While our findings provide insight into the range of AI professionals' views toward DOD and the factors underlying those views, they are preliminary and require cautious interpretation. Despite limited evidence of nonresponse bias based on company size or work

location, the survey had a low response rate and relatively high opt-out rate.[48] Additionally, we surveyed a small set of professionals with defined AI-relevant skills working in specific U.S. cities. Given the subject of the survey, professionals' response propensity may be shaped by pre-existing attitudes that we did not capture in the survey or control for in our analysis. For these reasons, additional research is needed to draw more definitive and generalizable conclusions on this topic. Continued survey research of relevant populations should dig deeper into the state of tech industry-DOD relations and the factors underlying different views toward collaboration. Meanwhile, ongoing CSET research explores private sector AI investment and innovation, as well as government access to AI talent and resources.[49]

Acknowledgements

For invaluable guidance and support over the course of this research, the authors thank Igor Mikolic-Torreira. The authors also thank Dewey Murdick and Melissa Flagg for their thoughtful input and guidance. For data assistance and review, the authors thank Rebecca Gelles, James Dunham, and Ilya Rahkovsky. For their valuable editing feedback, the authors would also like to thank Shelton Fitch and Danny Hague. For feedback on the survey instrument, the authors thank Amina Al Sherif, Andrew Kim, Eric Landree, James Ryseff, Tim Rudner, and Jennifer Melot. For feedback on this report, the authors thank Lt. Gen. Jack Shanahan, Robert Cardillo, Jenn Gamble, Miles Brundage, Helen Toner, and Max Langenkamp. The authors also thank focus group participants in San Francisco and Boston for their candid insights and helpful suggestions.

This survey was approved by the Georgetown University Institutional Review Board (Study #00001878).

Document Identifier: doi: 10.51593/20190050

Appendix

A. Sampling Methodology

Our sample of AI professionals was identified using stratified random sampling. First, we pulled individuals that report at least one of the following AI-relevant skills in their LinkedIn profile: machine learning, deep learning, artificial intelligence, computer vision, natural language processing, neural networks, and artificial neural networks. Then we grouped those individuals by city to identify the U.S. metropolitan areas with the highest number of AI-skilled employees (AI hubs). The top six metropolitan areas in terms of quantity of AI-skilled employees were San Francisco Bay Area, Greater Seattle, Los Angeles Metropolitan Area, New York City Metropolitan Area, Washington DC-Baltimore Metropolitan Area, and Greater Boston. Next, we identified the companies based in each city with the highest density of AI-skilled employees within various size bracket. To ensure we surveyed professionals from a range of companies, we divided those companies into three size brackets: 10-500 employees, 501-10k employees, and more than 10k employees. See the full list of companies included in the sample here. We then randomly selected AI-skilled employees within each company.

B. Survey Questionnaire

The full questionnaire presented to respondents in our main survey is available here.

C. Survey Distribution

To maximize response rates, we tested several pilot versions of the survey. The first pilot was distributed in March 2020 to 300 AI professionals and had a five percent response rate. This pilot included two follow-up questions that were removed from all subsequent versions of the survey and led to a minor revision of one question that asked whether respondents feel negative, neutral, or positive about working on DOD-funded AI projects. In two additional rounds of pilot surveys, each sent to a random sample of 300 AI professionals, we tested a shortened survey with only five questions and the full survey with revised survey invitations. The survey invitations varied by listing Georgetown University or CSET as the survey sponsor, explaining the research aims as related to national security or the AI workforce, and including short bios of research team members. We also sent a pre-notification email to 100 professionals notifying them of the research and

upcoming survey invitation. These pilots were distributed in March and April 2020 with response rates ranging from three to six percent. Based on response rates in the pilot distributions, the main survey, distributed in May 2020, included five core questions and an optional extension with 20 questions. The main survey had a response rate of three percent. Reported results include responses from each survey distribution, for a total of 160 responses, or a four percent response rate.

D. Sample Representativeness

Our sampling frame included people with AI-relevant skills who work in a U.S. AI hub at an AI company. We restricted our sampling frame to AI professionals working at AI companies in AI hubs, meaning metropolitan areas with the highest number of AI professionals, because we were interested in geographical variation in attitudes, which required sufficient responses from each region for meaningful comparisons. We collected estimates for the proportion of AI industry professionals in the U.S. working in AI hubs at AI companies using LinkedIn. We did not calculate estimates for the U.S. AI industry beyond these metropolitan areas or companies.

Table A displays the number of professionals that we identified (actual), sent the survey to (target), and that took the survey (respondents) for each company size and AI hub. For data availability reasons, we were not able to collect actual estimates from many small companies (less than 500 employees). To address this limitation, we intentionally over-sampled from larger companies and under-sampled from smaller companies, and calculated a target proportion of professionals that varies slightly from the actual proportion calculated based on LinkedIn data. This accounts for the discrepancy between the actual and target percentages reported in Table A. Therefore, we compare respondent proportions to our target proportions to assess representativeness, which indicates limited differences in the sample and population with respect to location and company size.

Table A. AI Professionals by AI Hub and Company Size

AI Hub		Company Size			
		10-500	501-10k	>10k	Total
Boston	Actual	10,100 (3%)	11,100 (4%)	9,600 (3%)	30,800 (10%)
	Target	69 (2%)	100 (3%)	252 (6%)	421 (11%)
	Respondents	5 (3%)	7 (4%)	8 (5%)	20 (13%)
Seattle	Actual	6,100 (2%)	5,800 (2%)	26,000 (9%)	37,900 (13%)
	Target	23 (1%)	68 (2%)	600 (15%)	691 (17%)
	Respondents	0 (0%)	5 (3%)	21 (13%)	26 (16%)
Los Angeles	Actual	10,300 (3%)	7,300 (2%)	8,400 (3%)	26,000 (9%)
	Target	33 (1%)	88 (2%)	187 (5%)	308 (8%)
	Respondents	1 (1%)	4 (3%)	8 (5%)	13 (8%)
New York City	Actual	20,100 (7%)	15,900 (5%)	26,000 (9%)	62,000 (20%)
	Target	61 (2%)	179 (4%)	564 (14%)	804 (20%)
	Respondents	2 (1%)	6 (4%)	22 (14%)	30 (19%)
San Francisco	Actual	33,800 (11%)	27,200 (9%)	58,000 (19%)	119,000 (39%)
	Target	113 (3%)	365 (9%)	959 (24%)	1437 (36%)
	Respondents	5 (3%)	20 (13%)	26 (16%)	51 (32%)
Washington D.C.	Actual	9100 (3%)	8600 (3%)	10,000 (3%)	27,700 (9%)
	Target	27 (1%)	79 (2%)	228 (6%)	334 (8%)
	Respondents	3 (2%)	3 (2%)	14 (9%)	20 (13%)
Total	Actual	89,500 (30%)	75,900 (25%)	138,000 (46%)	303,400
	Target	326 (8%)	879 (22%)	2790 (70%)	3995
	Respondents	16 (10%)	45 (28%)	99 (62%)	160

Source: LinkedIn Recruiter. Data pulled September 21, 2020.

E. Reasons for Not Taking the Survey

In each pilot distribution, we asked professionals who declined to participate to provide their reasons for not taking the survey. While a small number of nonrespondents provided at least one reason, reasons for nonresponse vary and are not strictly due to the survey topic, sponsor, length, or mode. Reasons were split between receiving the survey in error, having no opinions on the topic, not feeling comfortable, and not having the time. Several nonrespondents contacted the authors by email with their reason for not taking the survey, those reasons are included in Table B.

Table B. Reasons for Not Taking Survey

Reason	Frequency
Was sent this survey in error	5
Not comfortable taking this survey	5
Have no opinions on the topic	5
Don't have time	4
Need more information on the topic	3
Don't take surveys I get via email	3
Don't trust the survey sponsor	1
N	26

F. Other Benefits, Downsides, and Factors that Increase Willingness

The following are the full-text responses entered by respondents who selected "other" as a most compelling benefit or downside to working on DOD-funded AI projects or as a factor that would increase willingness to work on a DOD-funded AI project.

Other most compelling benefits:

- In my current position I do need funding, but in general DOD is one of the major contributing for technological advancement in the country. (p.s. I think Global Warming as a national threat should also be funded).

Other most compelling downsides:

- Fundamental lack of understanding by DOD officials about what AI can and cannot do.
- Some work, such as surveillance, used by irresponsible management can be used to limit civil privacy. The task should be designed such that it will get DOD needs with thought of the future too.
- I will only work on open source projects.
- Trade-off between maintaining secrecy vs whom to involve as collaborative partners (domestic and foreign)
- Losing control over the direction of the project.
- I am not sure if you state that the topics above are real or not. Since I do not know that I would stick with none.
- pays less

Other factors that increase willingness:

- International researchers being allowed access to projects
- More trustworthy administration that has transparency
- Ability to apply to commercial application
- transparent applications that align 100% with my personal morals, so very hard to quantify, but "liberal-leaning" things like combat climate change, combating domestic/systemic inequality, etc
- It involves saving lives and bringing green technology to fruition
- Lack of security clearance requirements
- Never Become Public

G. Additional Figures

Figure A. Willingness to Work on DOD-Funded AI Projects, Main Survey Only

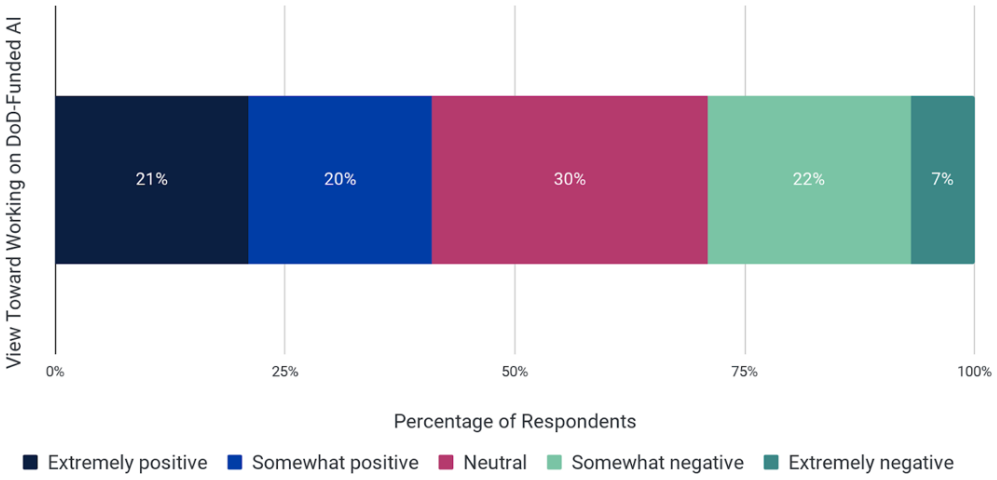

Question asked, "Do you generally feel negative, neutral, or positive about working on an AI project funded by a DOD grant or contract?" Respondents could choose from a 5-point scale; "extremely negative," "somewhat negative," "neither negative nor positive," "somewhat positive," or "extremely positive." Figure 1 in the report combines responses from both the pilot and main survey versions of the question. Respondents: 90. Proportions do not sum to 100 percent due to rounding. Source: CSET AI-DOD Survey 2020.

Figure B. Mean USG Agency Favorability by Feelings Toward Working on DOD-Funded AI Projects

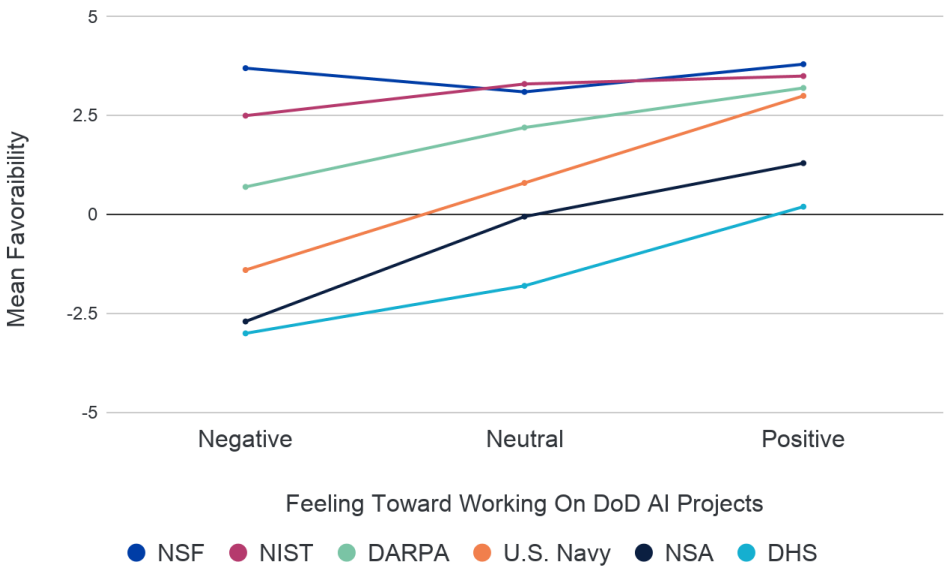

Difference in mean rating between respondents with negative and positive views toward working on DOD-funded AI projects was statistically significant at the 95 percent confidence level for the following agencies: DARPA, U.S. Navy, NSA, and DHS. A mean favorability between 0 and 5 indicates positive feelings and a mean favorability between 0 and -5 indicates negative feelings. A mean favorability of 0 indicates neutral feelings toward the agency. Respondents: 67. Source: CSET AI-DOD Survey 2020.

Endnotes

[1] DOD's AI Ethical Principles can be found at https://www.defense.gov/Newsroom/Releases/Release/Article/2091996/DOD-adopts-ethical-principles-for-artificial-intelligence/. While DOD's adoption of AI ethical principles was not solely for the purpose of reassuring the tech industry, it was likely a secondary aim. For example, shortly after the principles were announced in February 2020, former director of DOD's Joint AI Center, Lt. Gen. Jack Shanahan said, "If we had had the AI ethics principles three years ago [when launching Project Maven], and we were transparent about what we were trying to do and why we were trying to do it, maybe we would have a different outcome." Quoted in Sydney J. Freedberg Jr., "DOD Adopts AI Ethics Principles — But How Will They Be Enforced?" *Breaking Defense*, February 24, 2020, https://breakingdefense.com/2020/02/DOD-adopts-ai-ethics-principles-but-how-will-they-be-enforced/.

[2] Deputy Secretary of Defense, "Establishment of an Algorithmic Warfare Cross-Functional Team (Project Maven)," Department of Defense, April 26, 2017, https://www.govexec.com/media/gbc/docs/pdfs_edit/establishment_of_the_awcft_project_maven.pdf.

[3] The letter to Google CEO Sundar Pichai from signing employees can be found at https://static01.nyt.com/files/2018/technology/googleletter.pdf.

[4] Scott Shane and Daisuke Wakabayashi, "'The Business of War': Google Employees Protest Work for the Pentagon," *The New York Times*, April 4, 2018, https://www.nytimes.com/2018/04/04/technology/google-letter-ceo-pentagon-project.html.

5 Diane Greene, "Incorporating Google's AI Principles into Google Cloud," *Google Cloud (blog)*, June 7, 2018, https://www.blog.google/products/google-cloud/incorporating-googles-ai-principles-google-cloud/.

[6] Brad Smith, "Technology and the US military," *Microsoft On the Issues* (blog), October 26, 2018, https://blogs.microsoft.com/on-the-issues/2018/10/26/technology-and-the-us-military/.

[7] Employees of Microsoft, "An Open Letter to Microsoft: Don't Bid on the US Military's Project JEDI," *Medium*, October 12, 2018, https://medium.com/s/story/an-open-letter-to-microsoft-dont-bid-on-the-us-military-s-project-jedi-7279338b7132.

[8] For coverage of other cases of tech industry employee pushback against federal government contracts, see Kate Clark, "Protesters call on Salesforce to end contract with border patrol agency," *TechCrunch*, September 25, 2018, https://techcrunch.com/2018/09/25/protestors-call-on-salesforce-to-end-contract-with-border-patrol-agency/; Jamie Condliffe, "Amazon Is Latest Tech Giant to Face Staff

Backlash Over Government Work," *The New York Times*, June 22, 2018, https://www.nytimes.com/2018/06/22/business/dealbook/amazon-staff-facial-recognition-protest.html.

[9] Amy Zegart and Kevin Childs, "The Divide Between Silicon Valley and Washington Is a National-Security Threat," *The Atlantic*, December 13, 2018, https://www.theatlantic.com/ideas/archive/2018/12/growing-gulf-between-silicon-valley-and-washington/577963/.

[10] Scott Malcomson, "Why Silicon Valley Shouldn't Work with the Pentagon," *The New York Times*, April 19, 2018, https://www.nytimes.com/2018/04/19/opinion/silicon-valley-military-contract.html; Christopher M. Kirchhoff, "Why Silicon Valley Must Go to War," *The New York Times*, May 2, 2018, https://www.nytimes.com/2018/05/02/opinion/silicon-valley-pentagon.html.

[11] Eric Schmidt, "Eric Schmidt: I Used to Run Google. Silicon Valley Could Lose to China," *The New York Times*, February 27, 2020, https://www.nytimes.com/2020/02/27/opinion/eric-schmidt-ai-china.html; Amanda Macias, "Bezos says 'the country is in trouble' if big tech turns its back on the Pentagon: 'We are the good guys'," *CNBC*, December 7, 2019, https://www.cnbc.com/2019/12/07/bezos-says-country-in-trouble-if-big-tech-turns-its-back-on-the-pentagon.html. For other examples of public figures speaking out in favor of a collaborative relationship between the tech industry and DOD, see Bradley Peniston, "Google's Withdrawal from Pentagon AI Project Risks US Lives, Says Work," *Defense One*, June 26, 2018, https://www.defenseone.com/technology/2018/06/googles-withdrawal-pentagon-ai-project-risks-us-lives-bob-work/149280/; Michael R. Bloomberg, "Google Walks Away From America's Security, *Bloomberg*, June 6, 2018, https://www.bloomberg.com/opinion/articles/2018-06-06/google-s-decision-to-ditch-project-maven-is-a-grave-error; Ash Carter, "The morality of defending America: A letter to a young Googler," *The Boston Globe*, June 6, 2019, https://www.bostonglobe.com/opinion/2019/06/06/the-morality-defending-america-letter-young-googler/DTUQr3xXcP1svuGfYElNzI/story.html; Rachel Sandler, "Peter Thiel Says CIA Should Investigate Google For Being 'Treasonous'," *Forbes*, July 15, 2019, https://www.forbes.com/sites/rachelsandler/2019/07/15/peter-thiel-says-cia-should-investigate-google-for-being-treasonous/.

[12] Full-text of letter available at https://www.sec.gov/Archives/edgar/data/1321655/000119312520230013/d904406ds1.htm.

[13] Electronic Frontier Foundation (@EFF), "In a big win for ethical AI principles, Google will back away from military AI contracting. Congratulations to the Googlers and others who have worked hard to persuade the company to cancel its work on Project Maven." Twitter, June 1, 2018, https://twitter.com/eff/status/1002639117845295105; Will Knight, "One of the fathers of AI is worried about its future," *MIT Technology Review,* November 17,

2018, https://www.technologyreview.com/2018/11/17/66372/one-of-the-fathers-of-ai-is-worried-about-its-future.

[14] Tajha Chappellet-Lanier, "DIU says the relationship between DOD and Silicon Valley is just fine, thank you," *FedScoop*, October 30, 2018, https://www.fedscoop.com/diu-says-relationship-DOD-silicon-valley-just-fine-thank/; Aaron Mehta, "Former Symantec boss takes over the Defense Innovation Unit," *Defense News*, September 24, 2018, https://www.defensenews.com/pentagon/2018/09/24/former-symantec-boss-takes-over-the-defense-innovation-unit/.

[15] Loren DeJonge Schulman, Alexandra Sander, and Madeline Christian, "The Rocky Relationship Between Washington and Silicon Valley," (CNAS & The Copia Institute, July 2017), https://copia.is/wp-content/uploads/2017/07/COPIA-CNAS-Rocky-Relationship-Between-Washington-And-Silicon-Valley.pdf. In this study, nearly 80 percent of participants rated the state of collaboration between the communities as poor or very poor.

[16] Joseph Bernstein, "Survey: 51% Of Tech Industry Workers Believe President Trump Has A Point About the Media Creating Fake News," *BuzzFeed News*, February 23, 2019, https://www.buzzfeednews.com/article/josephbernstein/tech-industry-survey. In this survey, 59 percent of respondents agreed with the statement that "tech companies should work with the US government on military projects." Respondents were mostly professionals in IT services and technology roles. There is limited information provided on the survey sampling procedures, methodology, and question wording.

[17] Jack Poulson, "Reports of a Silicon Valley/Military Divide Have Been Greatly Exaggerated," *Tech Inquiry*, July 7, 2020, https://techinquiry.org/SiliconValley-Military/.

[18] Reported skills that we considered AI-relevant were: machine learning, deep learning, artificial intelligence, computer vision, natural language processing, neural networks, and artificial neural networks. To be included in the sample, one or more of these skills had to be listed as a "Skill" on the individual's LinkedIn profile. This sampling frame captured a wide-range of AI professionals in terms of technical expertise. To estimate the specific professional roles, we asked respondents to report their primary professional activities and results are reported in the findings section. The definition of DOD provided in the survey was "a U.S. federal agency that describes its mission as providing 'the military forces needed to deter war and ensure our nation's security.' DOD contracts or grants can support combat and non-combat applications."

[19] Our sampling procedure ensured variation in company size but due to data availability limitations, we over-sampled professionals working at large companies. Our sampling procedure also captured a range of AI companies in terms of AI sophistication and development. All companies included in our sampling frame are included in the appendix. The metropolitan areas included in our sampling frame are those with the highest number of workers with AI skills at the time of fielding the survey: San Francisco, New York City, Seattle, Boston, Washington, D.C.-Baltimore, and Los Angeles. We considered including other metropolitan areas with a growing AI-skilled workforce, such as Austin, Chicago, and St. Louis but decided against their inclusion. Therefore, we do not claim our findings are

generalizable beyond included metropolitan areas. See Appendix for more details on our sampling procedure and sample representativeness.

[20] Identifying AI professionals through LinkedIn could introduce some bias, if specific companies discourage employees from having profiles or employees from specific geographical locations use LinkedIn less often. While this is a potential source of bias, we believe it is minimal. The authors are not aware of any AI companies that explicitly discourage LinkedIn profiles among employees or specific areas in the United States with notably limited LinkedIn activity.

[21] The San Francisco focus group included six people and took place on November 6, 2019. The Boston focus group included nine people and took place on December 5, 2019. Each focus group lasted between 60 and 90 minutes and was led by a member of the research team and a recording of the focus group was transcribed by a member of the research team.

[22] The reported response rate was calculated according to American Association for Public Opinion Research's RR2: the number of completed and partial surveys divided by the number of eligible units in the sample, including cases of unknown eligibility (AAPOR, *Standard Definitions,* 2016). The minimum response rate (RR1), which does not count partial responses, is three percent. This is a low response rate, but not unexpected given our population of interest and survey topic. Additionally, our survey distribution coincided with the COVID-19 pandemic and subsequent transition to quarantine and remote work for many AI professionals. For purposes of comparison, the authors are not aware of any existing surveys that target this population about this topic and report response rates. Recent related surveys of AI talent have response rates ranging from 10 to 20 percent, but include AI professionals working in academia, who respond at higher rates. For example, a 2019 survey had a response rate of 17 percent among global AI/ML researchers, including 80 percent of respondents working in academia (see Baobao Zhang, Markus Anderljung, Lauren Kahn, Noemi Dreksler, Michael C. Horowitz, and Allan Dafoe, 2020, "Ethics and Governance of Artificial Intelligence: Evidence from a Survey of Machine Learning Researchers." Working paper). An earlier survey of a similar population had a 21 percent response rate, including 82 percent of respondents working in academia (see Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans, 2018, "When Will AI Exceed Human Performance? Evidence from AI Experts," *Journal of Artificial Intelligence Researcher*, https://arxiv.org/abs/1705.08807). A 2019 survey of U.S. AI PhDs had an 11 percent response rate, with 54 percent of respondents working in academia (see Catherine Aiken, James Dunham, and Remco Zwetsloot, "Career Preferences of AI Talent" (Center for Security and Emerging Technology, June 2020), cset.georgetown.edu/research/career-preferences-of-ai-talent/ and Catherine Aiken, James Dunham, and Remco Zwetsloot, "Immigration Pathways and Plans of AI Talent" (Center for Security and Emerging Technology, September 2020), cset.georgetown.edu/research/career-preferences-of-ai-talent/). A 2019 survey of AI researchers in the U.S. and UK had 184 respondents recruited through multiple international and university mailing lists, although no response rate is reported (see P.M. Krafft, Meg Young, Michael Katell, Karen Huang, and Ghislain Bugingo, "Defining AI in Policy versus

Practice," 2020 AAAI/ACM Conference on AI, Ethics, and Society, ssrn.com/abstract=3431304).

[23] The survey invitation that removed "defense" and "security" from the description of the survey topic had a slightly higher response rate than the original invitation but a slightly lower response rate than the invitation that requested participation in a shortened version of the survey, although each fell within the range of three to four percent.

[24] See Table B in Appendix for the reasons provided by nonrespondents for not taking the survey. Most indicated an unwillingness because they believe they were sent the survey in error, had no opinions on the topic, or did not have the time to complete the survey.

[25] Respondents could select up to three professional activities from the following: basic research, applied research, engineering, data analysis, QA testing, management or administration, professional service (e.g., consulting). Respondents were also offered the option to select "other" and enter additional professional activities. In pilot versions of the survey, we asked respondents, "Does your current professional position involve AI- or ML-related research or development?" Of pilot respondents, 81 percent said their current position involves AI- or ML-related research or development.

[26] There are no authoritative estimates of, and we were not able to estimate using our data, the proportion of AI professionals that are U.S. citizens or within each age bracket. We did not ask respondents to report their gender.

[27] A pilot version of the survey asked "Do you feel generally positive, neutral, or negative about: Accepting DOD grants for applied research? Accepting DOD grants for basic research? Working on DOD contracts at a company? Working for a company that has DOD contracts, independent of your involvement?" The question was revised to "Do you generally feel negative, neutral, or positive about working on an AI project funded by a DOD grant or contract?" in the main survey to reduce survey length. 42 respondents answered the pilot version of this question. For these respondents, their responses to "Do you feel generally positive, neutral, or negative about working on DOD contracts at a company?" were used to measure their general feelings toward working on a DOD-funded AI project reported in Figure 1.

[28] Although self-reported experience with an employer with a DOD contract is related to views toward working on a DOD AI project, we did not find that views differed based on whether the respondent is, based on federal contracting data, currently employed at a company that has had a DOD contract since 2015. This is likely because respondents who are/were aware of or involved in their employer's DOD contract-based work were more likely to report past/current experience and also to have that experience shape their general views, whereas respondents who are unaware or not involved in DOD contract work at their current employer were less likely to have it impact their views.

[29] Respondents were coded as East Coast if their work location was New York City, Boston, or Washington D.C. and West Coast if their work location was San Francisco, Seattle, or Los Angeles. Employer company size was broken down to greater or less than ten thousand

employees. Other subsample comparisons that resulted in no significant difference in positive or negative views toward working on DOD-funded AI projects include political interest and non-voting political activity, citizenship, and primary job activities. In terms of citizenship, we found a slightly higher proportion of U.S. citizens are positive while non-U.S. citizens tend to be neutral, but we cannot be confident in this comparison due to a very small number of non-U.S. citizen respondents. In terms of job activities, there was some evidence that respondents who said their job does not involve AI/ML research and development are more positive and that respondents whose primary professional activity is professional service (e.g., consulting) are more positive than researchers, engineers, and analysts.

[30] The average importance of funding source for negative respondents is higher (4.4) than among positive respondents (3.5) at the 95 percent confidence level. The difference translates roughly to a decrease from very important to moderately important.

[31] The U.S. government agencies that were ranked as more favorable by respondents with positive views of working on DOD-funded projects are National Security Agency (NSA), U.S. Navy, Defense Advanced Research Projects Agency (DARPA), and Department of Homeland Security (DHS). The difference in mean favorability ratings for these agencies between these groups is significant at the 95 percent confidence level. Mean ratings for the two other U.S. government agencies included in the survey, National Science Foundation (NSF) and National Institute of Standards and Technology (NIST) were not significantly different. See Figure B in Appendix.

[32] Respondents were randomly assigned to answer the benefits/downsides questions first. The questions asked, "People give different reasons for wanting to work on a DOD AI contract or grant. What reason do you find most convincing for working on a DOD AI contract or grant?" And "People give different reasons for not wanting to work on a DOD AI contract or grant. What reason do you find most convincing for not working on a DOD AI contract or grant?"

[33] Comparing responses from the open-text and multiple-choice questions, we find open-text responses focus more on doing good while multiple choice responses focus on more research-related considerations like interesting problems and resources. It may be that our multiple-choice options did not capture the breadth of reasons related to doing good. In particular, open-ended responses described the benefits of advancing scientific knowledge and conducting impactful research, which were not given as options in the multiple-choice question. Alternatively, it may be that respondents first think of a desire to do good, but when reminded of other potential benefits, such as interesting problems or research-related resources, they consider those more compelling.

[34] Several responses were coded as "other" because they were not similar to any other given reasons. For example, "Government intentions and politics." It is worth noting that only two responses cited disagreement with or concerns about working with DOD under the current administration.

[35] Mean willingness in humanitarian conditions was 0.9 points higher than mean willingness in battlefield conditions. This difference is statistically significant at the 99 percent confidence level.

[36] Mean willingness in the humanitarian conditions was 0.6 points higher than mean willingness in the department back-office condition. The difference is statistically significant at the 95 percent confidence level.

[37] This question was inspired by a question included in Zhang et al., "Ethics and Governance of Artificial Intelligence: Evidence from a Survey of Machine Learning Researchers." That survey asked AI/ML researchers who indicated opposition to researchers working on a military application of AI what action they would take if their organization decided to work on the application. Our question included the same set of actions but added expressing interest in working on the project and speaking out in favor of the project. Our question was also asked to all respondents, not just those who opposed the project.

[38] In Zhang et al., "Ethics and Governance of Artificial Intelligence," the authors also found avoiding working on the project to be the most common action surveyed AI/ML researchers would take in response to their organization working on an AI military application that they opposed, with expressing concern the second most common action considered.

[39] Question asked, "Have you ever worked or do you currently work for an organization that engages in DOD-funded research or contract-based work?" While all respondents currently work in industry, past employing organizations could include universities or public sector organizations.

[40] Combined, the 62 companies included in our sample have had $220 billion in DOD contracts since 2015. Data is from Bloomberg Government. Due to data limitations, we were unable to estimate the proportion of the population of AI professionals that have in the past or currently work for a company with DOD-funded contracts.

[41] In March 2019 a survey of UK tech professionals conducted by Doteveryone, a UK-based think tank, found that 63 percent want more opportunities to think about the impact of their products, 78 percent want more resources to help them assess the impact of their products, and 80 percent agreed that it is important to consider the potential consequences when designing new technologies. Agreement on these points was especially high among professionals with an AI/ML focus. See "People, Power, and Technology: The Tech Worker's View," (Doteveryone, May 2019), https://doteveryone.org.uk/wp-content/uploads/2019/04/PeoplePowerTech_Doteveryone_May2019.pdf.

[42] We did not find that agreement with the statement "My input matters for my organization's decisions about funding sources" was impacted by employer size or level of trust in their employer to develop AI in the best interest of the public.

[43] Zhang et al., "Ethics and Governance of Artificial Intelligence: Evidence from a Survey of Machine Learning Researchers." The authors also found the most trusted actors among

surveyed AI/ML researchers were "larger unaffiliated organizations" such as IGOs and international institutions, with the EU considered most trustworthy. Both survey samples had a similar level of trust in the U.S. government (a low level/not too much trust) and Chinese government (no trust) to develop AI. While our survey asked for the level of trust in U.S. tech companies, this survey asked respondents to rate ten specific companies. Facebook was considered not very trustworthy, but researchers' trust of Google, Microsoft, OpenAI, and DeepMind was similar to our results for trust in U.S. tech companies (mean trust of 1.4). We were not able to compare levels of trust in respondent's employer, academia, or the U.S. military because this survey did not include those actors.

[44] Baobao Zhang and Allan Dafoe, "Artificial Intelligence: American Attitudes and Trends" (Center for the Governance of AI, January 2019), https://governanceai.github.io/US-Public-Opinion-Report-Jan-2019/index.html.

[45] Zhang et al., "Ethics and Governance of Artificial Intelligence: Evidence from a Survey of Machine Learning Researchers" also compared results from their survey of AI/ML researchers to these results of the survey of the American public. They found similar differences, with the public placing more trust in the U.S. military to develop AI than researchers place in the militaries of the country where they work and researchers placing more trust in international institutions like the EU than the American public.

[46] A 2018 nationally representative survey found 47 percent of Americans are interested in politics most of the time and 22 percent donated to a campaign in the last year. See Stephen Ansolabehere, Brian F. Schaffner, and Sam Luks, "Cooperative Congressional Election Study 2018: Common Content. Release 2," August 28, 2019, Harvard University, http://cces.gov.harvard.edu. A 2019 Pew survey found 17 percent of Americans trust the government in Washington always or most of the time. See "Public Trust in Government: 1958-2019," (Pew Research Center, April 2019) https://www.pewresearch.org/politics/2019/04/11/public-trust-in-government-1958-2019/.

[47] Quote from letter available at https://www.sec.gov/Archives/edgar/data/1321655/000119312520230013/d904406ds1.htm.

[48] Seven percent of all survey contacts were opt-outs.

[49] Aiken et al., "Career Preferences of AI Talent"; Dakota Foster and Zachary Arnold, "Antitrust and Artificial Intelligence: How Breaking Up Big Tech Could Affect Pentagon's Access to AI," (Center for Security and Emerging Technology, May 2020), https://cset.georgetown.edu/research/antitrust-and-artificial-intelligence-how-breaking-up-big-tech-could-affect-pentagons-access-to-ai/; Melissa Flagg, "AI Hubs in the United States," (Center for Security and Emerging Technology, May 2020), https://cset.georgetown.edu/research/ai-hubs-in-the-united-states/; Zachary Arnold, Ilya Rahkovsky and Tina Huang, "Tracking AI Investment," (Center for Security and Emerging Technology, September 2020), https://cset.georgetown.edu/research/tracking-ai-investment/; Margarita Konaev, "U.S. Military Investments in Autonomy and AI," (Center for

Security and Emerging Technology, October 2020), https://cset.georgetown.edu/research/u-s-military-investments-in-autonomy-and-ai-executive-summary/.