

# Deepfakes and Scientific Knowledge Dissemination

**Christopher Doss** (✉ [cdoss@rand.org](mailto:cdoss@rand.org))

RAND Corporation

**Jared Monschein**

RAND Corporation

**Dule Shu**

Carnegie Mellon University

**Tal Wolfson**

Frederick S. Pardee RAND Graduate School

**Denise Kopecky**

Challenger Center for Space Science Education

**Valerie A. Fitton-Kane**

Challenger Center for Space Science Education

**Lance Bush**

Challenger Center for Space Science Education

**Conrad Tucker**

Carnegie Mellon University

---

## Research Article

### Keywords:

**Posted Date:** March 10th, 2022

**DOI:** <https://doi.org/10.21203/rs.3.rs-1408525/v1>

**License:** © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

**Additional Declarations:** No competing interests reported.

---

**Version of Record:** A version of this preprint was published at Scientific Reports on August 18th, 2023. See the published version at <https://doi.org/10.1038/s41598-023-39944-3>.

# Abstract

Science misinformation on topics ranging from climate change to vaccines have significant public policy repercussions. Artificial intelligence-based methods of altering videos and photos (deepfakes) lower the barriers to the mass creation and dissemination of realistic, manipulated digital content. The risk of exposure to deepfakes among education stakeholders has increased as learners and educators rely on videos to obtain and share information. We detail the results of the first study to understand the vulnerabilities of education stakeholders to science deepfakes and the characteristics that moderate vulnerability. We ground our study in climate change and survey individuals from five populations spanning students, educators, and the adult public. Our sample is nationally representative of three populations. We found that 33% to 50% of individuals cannot distinguish between authentic videos and deepfakes. All populations exhibit vulnerability to deepfakes, but vulnerability varies with age, political orientation, and trust in information sources. Adults and educators exhibit greater vulnerability compared to students, indicating that those providing education are especially susceptible. Vulnerability increases with more exposure to potential deepfakes, suggesting that deepfakes become more pernicious without interventions. Our results suggest that focusing on the social context in which deepfakes reside is one promising strategy for combatting deepfakes.

## 1. Introduction

Widespread use of the internet by learners of all ages has democratized the development and accessibility of educational materials<sup>1</sup>. The COVID-19 pandemic further solidified digital communication as a primary medium for information exchange among educators and learners<sup>2</sup>. Indeed, K-12 students are digital natives who use various online platforms such as YouTube to complete academic assignments<sup>3</sup>, with varying degrees of judgement for the reliability of the sources<sup>4</sup>. The amount of time children eight years old or less spent on YouTube has doubled between 2017 and 2020, portending an increased reliance on online media for even the youngest learners<sup>5</sup>. Social media and other online information sources have also facilitated lifelong education for continued professional and personal development<sup>6</sup>. However, the proliferation of misinformation on social media and other platforms has raised the risk of exposure to deliberately misleading educational content<sup>7,8</sup>.

The consequences of exposure to online misinformation range in severity and scale, often depending on context<sup>9</sup>. While many types of misinformation exist outside of the mainstream<sup>10</sup>, misinformation on science, technology, engineering, and math (STEM) related topics such as climate change and vaccines have had major public policy repercussions<sup>11-14</sup>. Social media and modern communication methods facilitate the rapid dissemination of misinformation, amplifying these impacts<sup>15,16</sup>. Misinformation campaigns tend to rely on undermining the consensus, highlighting uncertainty, undermining the credibility of leading figures and institutions, and disseminating pseudoscientific alternatives<sup>17</sup>.

The availability of open-source artificial intelligence (AI) algorithms has significantly lowered the barriers to altering videos and images in order to produce highly realistic, manipulated digital content (e.g., deepfakes)<sup>18-20</sup>. Generative neural networks (GNNs) are a class of deep neural network models that represent the state-of-the-art technique that can be leveraged to democratize the mass synthesis of manipulated digital content<sup>21,22</sup>. They have been used to fabricate images by training them to encode human features<sup>23</sup>, to manipulate images via replacing specific components of a digital image or video<sup>24</sup>, and to create videos via animation of a still image with the characteristics of a source video<sup>25</sup>.

In this study, we investigate the vulnerabilities of K-12 students, higher education students, teachers, principals, and general adult learners to deepfakes related to climate change and investigate potential population and video characteristics that can be leveraged in mitigation approaches. To date, the anticipated prevalence of deepfakes across societal contexts has motivated a large body of work seeking to develop algorithmic techniques to detect deepfakes<sup>26-37</sup>. However, these algorithms exhibit low rates of successful detection and are not robust across deepfake types, content format, content characteristics, and datasets<sup>20,38</sup>. Parallel efforts to advance user-focused solutions are nascent and characterized by high failure rates<sup>39</sup>. Recent work indicates that human-machine teams show promise for overcoming these challenges via complementary approaches to identifying deepfakes<sup>40-42</sup>; these studies, however, don't account for the social and individual characteristics that modulate individuals' vulnerabilities to deepfakes. The successful design, development, and deployment of human-machine teams for deepfake detection and other mitigation strategies requires a comprehensive understanding of individuals' abilities to successfully detect deepfakes, personal characteristics that moderate individuals' vulnerability to deepfakes, and digital content characteristics that influence successful detection<sup>43,44</sup>. The enabling data, however, have not yet been made available.

The detection and mitigation against deepfakes are particularly needed within STEM education given increasing access to and reliance on readily available digital educational content by both youth and adult learners<sup>45</sup>. To date, work investigating the vulnerabilities of K-12 students to STEM misinformation has tended to focus on deliberately falsified text-based content and media literacy<sup>46-50</sup>. A limited number of studies have investigated adults' vulnerability to deepfakes, but these have been limited to deepfakes depicting politicians, how this exposure impacts voters' attitudes toward politicians depicted in the videos, and how vulnerability can be moderated by personal characteristics (e.g., religious convictions, political orientation), as well as attempted inoculations within these contexts<sup>41,51,52</sup>.

Climate change is a particularly compelling aspect of STEM to explore because the polarized nature of climate change has left this domain particularly vulnerable to digital misinformation. Climate change misinformation outside of deepfakes is pervasive and typically relies on recipients'

motivated cognition to protect against ideologically or economically threatening scientific evidence<sup>53, 54</sup> to gain traction. Weak media literacy skills, particularly among K-12 students, has also been shown to moderate susceptibility<sup>55</sup>. Historically, producing convincing fabricated or manipulated digital content (data, videos, audio, etc.) related to climate change has been much more challenging<sup>54</sup>. However, the emergence of AI algorithms as a tool for manipulating digital content – particularly to create deepfakes – increases the risk of exposure to convincing climate change misinformation<sup>56</sup>. Additionally, it's currently unknown if deepfakes present novel threat vectors that can take advantage of similar vulnerabilities as mainstream climate change misinformation or if deepfakes expand the misinformation attack surface.

To investigate the vulnerabilities of the education system climate change deepfakes, we fielded a survey that embedded a series of randomly assigned authentic or deepfake videos on climate change. We then asked respondents to identify the video as authentic or manipulated and gathered information regarding respondents' demographics, background knowledge of climate change, learning habits, and perspectives on deepfakes. We found that between one third to over half of survey responses were unable to correctly identify the authenticity of videos *regardless* of whether the video was authentic or a deepfake. In aggregate, U.S. adults and educators were less likely to correctly identify deepfake videos than authentic videos, while middle school and higher education students were more likely to identify deepfake videos than authentic videos. However, vulnerability fluctuates across individual deepfake videos and can be quite severe. Heterogeneity analyses indicated that an individual's susceptibility varies as a function of age, political orientation, and trust in information sources. Further, vulnerabilities increased dramatically as individuals were exposed to more potential deepfakes, suggesting that deepfakes can become more pernicious without educational interventions. An analysis of video characteristics that respondents reported drove their decisions indicated that the social context in which deepfakes are embedded could provide a promising approach for educational mitigation strategies. We conclude by discussing the implications of these results on the development of technical and social mitigation strategies for combatting STEM-focused deepfakes.

## 2. Survey Instrument And Sample

We fielded our survey to five key populations: nationally representative samples of (1) adults in the U.S. 18 years of age and older, (2) U.S. K-12 teachers, and (3) U.S. K-12 principals, (4) a large sample of middle school students from three states across the U.S., and (5) a sample of undergraduate and graduate students at Carnegie Mellon University (CMU). See methods section for more detail on these samples.

Our main research question, which sought to understand the vulnerability of each population to the deepfake videos, is:

1. What is the effect of receiving a deepfake video on an individual's ability to correctly identify a video's authenticity and how does that effect vary by population?

Our secondary research questions, which leverage the contextual questions in the surveys to uncover drivers of vulnerabilities and potential avenues for deepfake mitigation, were as follows:

1. What characteristics of the video do respondents analyze to determine the authenticity of the video and how are they related to the probability of correctly identifying the video's authenticity?
2. How does the effect of a deepfake video on an individual's ability to correctly identify a video's authenticity vary by respondent background characteristics and beliefs?

The survey instrument took about ten minutes to complete and was divided into two portions. One portion presented the respondent with four videos, each about 10–15 seconds in length. Each video featured one of four speakers: (1) Timothy Gallaudet, an oceanographer and former Acting Administrator of the National Oceanic and Atmospheric Administration; (2) Richard Lindzen, Professor Emeritus of Meteorology at the Massachusetts Institute of Technology; (3) Greta Thunberg, noted climate change activist; and (4) Naomi Seibt, anti-climate change activist. These speakers were chosen to expose respondents to a variety of views and types of speakers on climate change that youth and adult learners are likely to encounter. Gallaudet and Lindzen are credentialed speakers with positions of scientific authority with opposing views on climate change. While Gallaudet is a climate change believer, Lindzen is a climate change skeptic. Greta Thunberg and Naomi Seibt are both younger activists and represent the climate change and anti-climate change perspective, respectively.

The four videos were drawn from a bank of eight possible videos. For each speaker, the bank contained an authentic video in which the speaker espoused their view of climate change and an AI generated deepfake video where the speaker is made to espouse their opposite view of climate change. Thus, climate change believers are made to be climate change skeptics and vice versa. Each respondent received one video from each speaker and was randomly assigned to receive an authentic or deepfake video of each speaker. A participant could view an authentic video of one speaker, but a deepfake of another speaker. In contrast to prior studies<sup>41</sup>, we did not tell participants the probability of being exposed to a deepfake because stakeholders would not be privy to such information when encountering deepfakes in real-life. We randomized the order of the speakers for each respondent to avoid any effects of being exposed to prior authentic or deepfake videos.

The second portion of the survey elicited the views of the respondents on a variety of dimensions including: personal climate change beliefs, perceptions of the scientific consensus on the causes of climate change, knowledge of climate change, political affiliation (asked only of adults),

information sources respondents use to learn about climate change, trust in those information sources, internet use habits, use of social media platforms, perceptions of the risks associated with and the prevalence of deepfakes, and demographic questions.

### 3. Results

#### 3.1 The Ability of Respondents to Correctly Detect Deepfake and Authentic Videos

Across all videos, substantial percentages of respondents did not correctly identify a video's authenticity when receiving deepfake *or* authentic videos (Fig. 1). When receiving authentic videos, the percentage of correct responses ranged from 50 percent (middle school students) to 66 percent (CMU students). When receiving deepfake videos, the percentage of correct responses ranged from 46 percent (adults) to 80 percent (CMU students). Figure A1 in Appendix A shows that among the responses that were not correct, a substantial number of responses were either "Cannot Tell" (between 10 to 20 percent of responses) or incorrect (between 11 to 33 percent of responses). These results imply that the possibility of receiving a deepfake video introduces ambiguity with a third to over a half of education stakeholders potentially being unable to correctly identify the authenticity of a video. These vulnerabilities to climate change-related deepfakes are similar to the vulnerabilities of adults to politically-oriented deepfakes focused on elections and context-agnostic deepfakes<sup>41, 51, 52</sup>, indicating that vulnerabilities to deepfakes may be consistent across contexts. These rates also closely mirror the susceptibility rates of individuals to digital images altered manually without artificial intelligence<sup>41, 57</sup>.

Our results further indicate that across all videos, receiving a deepfake lowered the probability that adults and educators correctly identified the authenticity of the video by 6 percentage points (principals) to 11 percentage points (adults) compared to when they received an authentic video. Based on the percentage of respondents that correctly identified authentic videos, these estimates represent a 10 percent (principals) to 19 percent (adults) decrease in correct responses. In contrast, student populations were better able to detect deepfake videos (compared to authentic videos) with an 8 percentage point increase in correct responses among middle school students and a 14 percentage point increase in correct responses from CMU students. This translated to a 16 and 21 percent increase in correct responses, respectively. Regression models presented in Table A1 in Appendix A show that these differences were statistically significant to the one percent level. Further, the effects of receiving a deepfake on the probability of responding "Cannot Tell" were small and statistically insignificant, implying that the deepfake videos induced participants to respond incorrectly in the adult and educator populations or correctly in the student populations.

These overall results, however, hide important heterogeneity as respondents' accuracy in correctly identifying deepfake and authentic videos varied across individual videos (Figures A2 to A4 in Appendix A). CMU students were the only population more likely to correctly identify deepfake videos than the authentic videos across all speakers, likely because of their experience and expertise in artificial intelligence and machine learning. All other populations were less able to correctly identify at least two deepfake videos. Adults, teachers, and principals, and middle school students were less able to detect deepfake videos of Timothy Gallaudet and Richard Lindzen. The largest gap was observed for the U.S. adults that viewed the Richard Lindzen video. Adults were 39 percentage points less likely to have correctly identified the deepfake video compared to the authentic video. This gap is a nearly a 50 percent decline in the proportion of correct responses compared to the authentic video. Results for the Greta Thunberg and Naomi Seibt videos were mixed, with middle school students more likely to detect deepfakes of both speakers while adults, teachers, and principals were equally or less able to detect the deepfakes of these speakers. Table A2 in Appendix A shows that these effects were highly significant in most cases. More research needs to be done to understand the drivers of this heterogeneity in effects, though overall our results show that no population in our study was immune to being fooled by deepfake videos.

#### 3.2 Video Characteristics Moderating Respondents' Vulnerability to Deepfakes

Identifying the characteristics of deepfake videos that moderate an individuals' vulnerability is an important step towards developing countermeasures. Understanding which aspects of a video respondents used in their decision-making process and how that differed when receiving an authentic or manipulated video provides an understanding of the respondents' cognitive processes. For example, if receiving a fake video resulted in greater (or less) use of technical aspects of the video such as facial features or overall quality, then respondents are making decisions in part based on overall or specific flaws in the video. Similarly, if receiving a fake video resulted in a greater (or less) use of social aspects of the video such as familiarity with the speaker's views or overall credibility of the content, then respondents are using contextual knowledge in their decision-making process. We then explore whether leveraging these aspects of videos are associated with the correct identification of authentic and manipulated videos.

The results in Table 1 indicate that respondents tended to use visual aspects of deepfake videos (e.g., facial features, video background, overall video quality) more often and social aspects of the video (e.g., credibility of the information being shared) or the audio quality less often when presented a deepfake. Principals, teachers, and middle school students were statistically significantly more likely to assess the overall quality of deepfake videos than authentic videos, while both teachers and CMU students were more likely to examine the facial features presented in deepfakes.

**Table 1:** Relationship Between Reported Aspect of Video Analyzed and Receiving Deepfake Video

	Adults	Principals	Teachers	Middle School Students	CMU Students
Familiar With Person's Views	-0.036 (0.024)	-0.027+ (0.015)	-0.027+ (0.015)	-0.017 (0.013)	0.036 (0.039)
Video Quality	0.051 (0.034)	0.086** (0.023)	0.074** (0.023)	0.047* (0.020)	0.034 (0.062)
Background	0.008 (0.024)	0.002 (0.020)	-0.011 (0.020)	-0.017 (0.019)	0.048 (0.055)
Facial Features	0.011 (0.035)	0.009 (0.021)	0.048* (0.022)	0.035+ (0.019)	0.115* (0.047)
Audio	-0.047 (0.033)	-0.039* (0.020)	-0.087** (0.020)	-0.045* (0.020)	-0.185* (0.072)
Credibility of Content	-0.059* (0.026)	-0.101** (0.022)	-0.065** (0.022)	-0.048** (0.017)	-0.017 (0.058)
N(Respondent-by-Video)	3044	2,960	2,568	3,220	348
N(Respondent)	761	740	642	805	87
Notes: Each cell presents the results of a separate regression of whether a respondent reported analyzing an aspect of the video and receiving a fake video. All models include speaker, video order, and respondent fixed effects. Regressions models used to analyze adult, principal, and teacher panels are weighted to retain national representation. Standard errors are clustered by respondent. + indicates $p < 0.10$ ; * $p < 0.05$ ; ** $p < 0.01$ .					

The results in Table 2 indicate the relationship between using these video aspects in the decision-making process on the likelihood of correctly identifying a video's authenticity. These results indicate using a videos' overall quality and the speaker's facial features in the decision-making process is associated with a significantly increased probability of correctly identifying deepfake videos, but not authentic videos. Adults, principals, and teachers were 20–25 percentage points more likely to accurately identify deepfakes when assessing the video quality while principals, teachers, and middle school students were 13–20 percentage points more likely to correctly identify deepfakes when analyzing facial features. Analyzing the audio quality of the videos was not associated with higher rates of correctly identifying deepfakes videos but did lead to 13–22 percentage point improvements in correctly identifying authentic videos by adults, principals, and middle school students. Thus, the usage of visual aspects of a video in the decision-making process when receiving a deepfake reduced viewers' vulnerability. This finding is aligned with previous work that similarly found that focusing on the visual aspects of videos can reduce the susceptibility of adults to deepfakes<sup>58</sup>.

**Table 2:** Relationship Between Reported Aspect of Video Analyzed and Correctly Identifying Video Authenticity, By Video Deepfake Status

	Adults		Principals		Teachers		Middle School Students		CMU Students	
	Real	Deep fake	Real	Deep fake	Real	Deep fake	Real	Deep fake	Real	Deep fake
Familiar With Person's Views	0.381** (0.089)	0.217+ (0.118)	0.479** (0.069)	0.223** (0.082)	0.389** (0.071)	0.101 (0.085)	0.409** (0.067)	-0.136 (0.084)	0.455* (0.199)	0.044 (0.192)
Video Quality	-0.128 (0.081)	0.251* (0.097)	-0.062 (0.049)	0.261** (0.058)	-0.134* (0.058)	0.210** (0.059)	0.020 (0.049)	0.085+ (0.052)	-0.059 (0.147)	-0.124 (0.139)
Background	-0.135 (0.125)	0.088 (0.102)	-0.075 (0.053)	0.083 (0.070)	0.010 (0.062)	-0.023 (0.076)	0.003 (0.046)	0.126* (0.055)	-0.135 (0.152)	-0.075 (0.156)
Facial Features	-0.033 (0.069)	0.152+ (0.082)	0.063 (0.047)	0.154** (0.057)	0.011 (0.058)	0.133* (0.061)	0.090 (0.056)	0.197** (0.055)	-0.011 (0.163)	-0.154 (0.134)
Audio	0.222** (0.076)	-0.001 (0.084)	0.222** (0.052)	-0.085 (0.058)	0.240** (0.060)	-0.102 (0.065)	0.133* (0.054)	-0.110* (0.051)	0.152 (0.163)	-0.074 (0.123)
Credibility of Content	0.157* (0.063)	-0.043 (0.074)	0.179** (0.050)	-0.073 (0.057)	0.064 (0.057)	-0.091 (0.056)	0.093+ (0.056)	-0.084 (0.067)	0.214 (0.131)	-0.065 (0.153)
N(Respondent-by-Video)	1,600	1,444	1,514	1,446	1,283	1,285	1,600	1,620	169	179
N(Respondent)	727	693	692	685	604	596	757	755	82	84

Notes: Each column presents the results of a separate regression of video authenticity on aspects respondents reported analyzing when making their decisions. All models include speaker, video order, and respondent fixed effects. Regressions models used to analyze adult, principal, and teacher panels are weighted to retain national representation. Standard errors are clustered by respondent. + indicates  $p < 0.10$ ; \*  $p < 0.05$ ; \*\*  $p < 0.01$ .

However, Table 2 also shows that propensity to use social aspects of the video less when receiving a deepfake video may not always be optimal, particularly for adults. Both adults and teachers that reported using their familiarity with the speaker's views in their decision-making process were more likely to correctly identify deepfakes by about 22 percentage points ( $p < 0.10$  for adults and  $p < 0.05$  for teachers). Using the credibility of the content also helped adults and teachers accurately identify authentic videos. Respondents who reported that they were familiar with the speaker's views were also statistically significantly more likely to correctly identify authentic videos in every population and with estimates that range from 38 to 48 percentage points. Educating information consumers to assess the social context that surround deepfakes such as content credibility and the speaker identity therefore present new focal points for strategies to mitigate vulnerabilities to deepfakes that may be more robust than a focus on technical aspects. Indeed, over the long term the value of assessing the technical aspects of deepfake videos is likely to decline as deepfake generation technologies are expected to continue their rapid advancement towards producing content that is indistinguishable from authentic videos<sup>59</sup>.

### 3.3 Respondent Characteristics Moderating Vulnerability to Deepfakes

The design of the survey instruments fielded in this work enabled an understanding of how respondent characteristics and beliefs moderated their vulnerabilities to deepfakes. Table 3 shows those characteristics that most consistently moderated affects across populations, while Table A5 in Appendix A shows that factors such as respondents' race, beliefs in climate change, perceived ability to detect deepfakes, perceived risk of deepfakes, climate change knowledge, frequency of consumption of information, frequency of social media use, urbanicity, working in or teaching a science or math related field, income, and education did not consistently moderate ability to detect video authenticity across populations.

**Table 3:** Moderation of Participant Background Characteristics on Probability of Correctly Identifying Deepfake Video

	Adults		Principals		Teachers		Middle School Student	CMU Students		
Panel A: Age In Years										
Age	-0.006*	-0.004*	-0.003	-0.003	-0.005*	-0.006**	-0.019	-0.035*	-0.029	-0.032
	(0.003)	(0.002)	(0.003)	(0.003)	(0.002)	(0.002)	(0.020)	(0.015)	(0.044)	(0.037)
Panel B: Political Orientation (Reference Category: Liberal)										
Conservative	-0.151+	-0.150*	-0.096	-0.097*	0.120+	0.132**	—	—	0.076	0.010
	(0.087)	(0.065)	(0.062)	(0.047)	(0.062)	(0.045)	—	—	(0.193)	(0.187)
Moderate	-0.077	-0.062	-0.126*	-0.116*	0.145*	0.166**	—	—	0.033	0.008
	(0.106)	(0.073)	(0.058)	(0.046)	(0.066)	(0.051)	—	—	(0.184)	(0.155)
Prefer Not To Say	-0.095	-0.080	0.125	0.081	0.021	0.077	—	—	-0.158	-0.220+
	(0.128)	(0.101)	(0.091)	(0.068)	(0.080)	(0.063)	—	—	(0.147)	(0.131)
Panel C: Trust in Information Sources										
Trust	-0.043	-0.022	-0.034	-0.046*	-0.043	-0.048*	-0.071**	-0.049**	-0.094+	-0.119*
	(0.029)	(0.024)	(0.027)	(0.021)	(0.026)	(0.020)	(0.024)	(0.019)	(0.055)	(0.054)
Panel D: Order of Videos Seen (Reference Category: First Video)										
Second Video	0.066	0.026	-0.121+	-0.112*	-0.096	-0.098+	-0.110+	-0.084	-0.265+	-0.315*
	(0.104)	(0.079)	(0.064)	(0.050)	(0.075)	(0.056)	(0.067)	(0.051)	(0.150)	(0.121)
Third Video	0.070	0.031	-0.170**	-0.143**	-0.192*	-0.215**	-0.129*	-0.179**	-0.357*	-0.472**
	(0.110)	(0.082)	(0.064)	(0.051)	(0.075)	(0.056)	(0.063)	(0.048)	(0.170)	(0.133)
Fourth Video	0.026	0.019	-0.149*	-0.138**	-0.126+	-0.175**	-0.222**	-0.223**	-0.257	-0.352**
	(0.110)	(0.084)	(0.064)	(0.051)	(0.070)	(0.053)	(0.063)	(0.048)	(0.155)	(0.114)
Respondent Fixed Effects	✓		✓		✓		✓		✓	
Respondent Controls		✓		✓		✓		✓		✓
N(Respondent-by-Video)	3044		2,960		2,568		3,220		348	
N(Respondent)	761		740		642		805		87	
Notes: Each cell presents the interaction term of a separate regression of whether a respondent correctly identified the authenticity of a video on the main effect for seeing a deepfake video and an interaction with the characteristic indicated by the row headers. All models include speaker and video order fixed effects. Models include respondent fixed effects or respondent controls as indicated in the table. In all models, respondent controls are panel specific variables listed in Table 4. The middle school and CMU student samples also include state fixed effects. Regressions models used to analyze adult, principal, and teacher panels are weighted to retain national representation. Standard errors are clustered by respondent. + indicates $p < 0.10$ ; * $p < 0.05$ ; ** $p < 0.01$ .										

In both the general adult population and the teacher population, the ability to detect deepfake videos declined with increasing age. We found that each year of age was associated with a 0.6 percentage point reduction in likelihood of correctly identifying deepfakes (Panel A of Table 3). Similar trends were observed for the principal, middle school, and CMU panels but these relationships were not as robustly statistically significant across specifications. This finding contrasts with earlier work that found vulnerability to digital misinformation decreased as a function of increasing age<sup>60-62</sup>. The digital misinformation in these studies did not include deepfakes, which indicates that deepfakes present a distinct modality of misinformation that affects information consumers differently than other misinformation pathways. These results suggest that older populations specifically should be prioritized as target audiences for deepfake mitigation strategies. This is the first study identifying higher levels of vulnerability to deepfakes among older information consumers.

The relationship between vulnerability to deepfakes and respondents' political orientation also indicates that deepfakes may be considered a distinct type of misinformation. Similar to the extant literature that indicates individuals who identify as politically conservative tend to be more vulnerable to climate change misinformation<sup>62, 63</sup>, we found that self-identified conservatives were less likely than self-identified liberals to correctly identify deepfakes by 10–15 percentage points in the adult and principal populations (Panel B of Table 3). However, among teachers, self-identified conservatives were better than liberals at identifying deepfakes by about 13 percentage points. The cause for this reduced vulnerability among self-identified politically conservative teachers is unknown. Future work is needed to further understand how participant beliefs and perspectives across populations intersect with the content and social context of deepfakes to produce these more nuanced results.

However, similar to other modalities of STEM misinformation<sup>64, 65</sup>, we found that an individual's trust in information sources (e.g., social media, newspapers, etc.) influenced their vulnerability to deepfakes. Panel C of Table 3 shows that for principals, teachers, middle school students, and CMU students, a one standard deviation increase in our factor of information trust (see Appendix C for more information on the factor) reduced the probability of correctly identifying a deepfake by about 5 to 12 percentage points. A similar relationship was observed among U.S. adults, although this was not statistically significant in any specification. This finding suggests that media literacy approaches, which aim to combat misinformation by teaching information verification techniques<sup>49, 55</sup>, may be a promising mitigation strategy to reduce vulnerability to deepfakes.

Our results also indicate that other mitigation strategies, such as training information consumers to identify deepfakes via repeated exposure, are unlikely to succeed and may worsen vulnerabilities (Panel D of Table 3). Among teachers, principals, middle school students, and CMU students, respondents were 13 to 47 percentage points less likely to correctly identify video authenticity for the third and fourth video consumed (no relationship between exposure and vulnerability was detected for U.S. adults). These results further indicate that without education on deepfakes, consistent exposure to potential deepfakes may cause more confusion that leads to less accuracy in video authenticity detection and a worsening effect of deepfake misinformation.

## 4. Conclusions

Society is at an inflection point regarding misinformation as social media platforms and the ubiquity of technology have allowed misinformation to proliferate and emerging technologies have reduced the barriers to production. STEM misinformation poses a substantial threat to society because it can impede policy reforms needed to combat global challenges such as climate change. To date, technological limitations meant that misinformation was usually disseminated in print, through the manipulation of still photos, or by selectively editing videos. Generative neural networks facilitate the creation of highly realistic deepfake videos that don't require specialized expertise to create, opening up a new front in the distribution of and fight against misinformation. Understanding vulnerabilities to deepfakes among key stakeholders and the factors that moderate susceptibility is critical for creating and implementing robust mitigation strategies. This study represents a first step at achieving those aims by analyzing a nationally representative set of U.S. adults, teachers, and principals, a large sample of middle school students throughout the country, and a sample of technically-oriented undergraduate and graduate students. Our results have three broad implications.

First, we found that deepfakes are already of sufficient quality as to introduce substantial confusion that leaves all education stakeholders vulnerable. Across all populations, between a third to over a half of respondents were unable to correctly identify a video regardless of its authenticity. Across all videos, adult populations were more vulnerable than CMU and K-12 students, though all populations except for CMU students were less able to correctly identify the deepfake version of the video (compared to the authentic version) for at least two of four speakers. This gap between correctly identifying deepfake videos versus authentic videos was seen to be as large as 39 percentage points. These results have severe implications for the vulnerability of the U.S. education system. Vulnerabilities of K-12 students open the possibility that students take science misinformation as fact and develop a flawed view of science, the scientific process, and the policy implications of science. If these flawed views go unchallenged, they can solidify as students mature and enter positions in society where they can more directly influence policy. The more severe vulnerability of adults, and particularly of teachers and principals, is perhaps more alarming. Teachers and principals must synthesize and communicate contemporary issues in science to students and vulnerabilities to misinformation can translate to inadvertent teaching of information to students. Further, educators are even less equipped to debunk any misinformation students bring into the classroom as students increasingly become exposed to and rely on online information. The vulnerability of adults means that parents may not be able to debunk misinformation their children are exposed to and may play a role in spreading it.

Second, we found that focusing on technical aspects such as overall video quality or facial features can help respondents more accurately detect deepfakes but analyzing the social aspects such as knowledge of the speaker or the credibility of the content can help more accurately identify both deepfake *and* authentic videos. Thus, while prior studies focused on the more technical aspects of videos and the interplay between human and algorithmic detections, our results show that the social context of videos is also critical to consider. Educating information consumers to assess social characteristics of digital content may therefore present new focal points for strategies to mitigate vulnerabilities to deepfakes that may be more robust than a focus on technical aspects. Over the long term, deepfake generation technologies are expected to continue their rapid advancement towards producing content that is indistinguishable from authentic videos, eroding the utility of visually assessing technical aspects of deepfake videos.

Finally, we found several personal-level characteristics of information consumers that can guide the development and implementation of mitigation strategies. First, older individuals tended to exhibit higher levels of vulnerability and so specific educational interventions should be tailored for this population. Second, higher trust in information sources, including in print, peer, and online sources, was found to raise respondents' vulnerability to deepfakes. This finding suggests that the traditional media literacy approaches may also be effective in combatting deepfakes. Our results also suggest that other types of mitigation strategies, such as those that rely on repeated exposure, are unlikely to succeed and may worsen vulnerabilities. These latter results further highlight the need for education on deepfake detection, as repeated exposure to potential deepfakes can increase confusion and allow deepfakes to become a more potent medium by which to spread misinformation.

## 5. Limitations

This is the first comprehensive analysis of societal vulnerabilities to deepfake videos. However, this study has some important limitations. First, the characteristics of videos between speakers were not varied in systematic ways, preventing us from understanding the discrete characteristics of deepfake videos that drive the heterogeneity in responses and vulnerabilities observed. Further research is needed to understand how specific features of deepfakes drive individuals' vulnerability. Second, self-reported information on sensitive questions such as climate change, political ideology, and deepfakes can suffer from social desirability bias and may not always reflect the true view of the respondent. Though randomization meant that our main vulnerability results were not affected by this bias, the relationships between deepfake vulnerability and these aspects may have been affected. Third, the analysis of the relationships between video aspects and the probability of respondents correctly identifying deepfake and authentic videos are the only correlational analyses in this study. Though these results point to promising avenues for deepfake mitigation, more research is needed to establish causal connections. Finally, this study analyzed the specific type of misinformation where people's views on climate change were manipulated. The results may not generalize to misinformation on other topics such as vaccinations and it may not generalize to deepfakes that manipulate content in other ways. Further research is needed to understand how these results generalize to other ways in which deepfake videos can be operationalized.

## 6. Methods

### 6.1. Additional Survey Details

The survey consisted of two parts. In one part, respondents were asked to watch four videos. After each video, respondents were asked about the authenticity of the video and which aspects of the video helped them make their decision. The respondent could have identified the video as "Definitely Fake," "Probably Fake," "Cannot Tell," "Probably Real," or "Definitely Real." These options allowed respondents to express a view on the authenticity of the video while providing an option (cannot tell) that suggests confusion in the respondent. Options regarding the aspects of the video that helped respondents make their decision include technical aspects (quality of the video, facial features of the speaker, background, and audio) as well as social aspects (familiarity of the respondent with the speaker's view and credibility of the content of the video). These characteristics were chosen based on the literature on the creation of deepfakes and the social science literature on misinformation<sup>20, 54</sup>.

The second part of the survey elicited the views of the respondents on a variety of dimensions. The constructs were chosen based on research that showed a relationship between these constructs and vulnerability to misinformation<sup>12, 21, 54, 55, 58, 63</sup>. Where possible, we used or modified questions and scales that have been used in previous misinformation studies. Where no questions could be found, we created our own items. Appendix B contains the survey instrument and lists the sources of questions used from previous studies.

Theoretically, viewing the climate change associated videos prior to answering questions eliciting views on climate and deepfake videos could affect a respondent's answer on those questions and vice versa. To guard against this possibility, we randomized whether the individual received the video or contextual questions first.

### 6.1. Deepfake Creation

The AI generated deepfake videos used in the survey were created via video-to-video synthesis<sup>66</sup>, a technique that generates artificial video content by transforming another video. The backbone of our video-to-video synthesis tool was the First Order Motion (FOM) model<sup>25</sup>, a machine learning model that takes two inputs, a driving video and a reference image, and yields a synthesized video that is an animation of the reference image according to the motion of objects in the driving video. To produce deepfakes related to climate change, we selected driving videos of individuals speaking in front of a camera about climate change and reference images of distinct individuals speaking in front of a camera. The audio component of the driving video was then added to the synthesized video. Following this procedure, we created four deepfake videos for our survey as shown in Figure A5 in Appendix A.

### 6.2 Study Populations

This survey leveraged three standing, nationally representative panels at the RAND Corporation: (1) The American Life Panel (ALP), a sample of adults 18 years of age or older; (2) the American School Leader Panel (ASLP), a sample of U.S. K-12 public school principals; and (3) American Teacher Panel (ATP), a sample of U.S. K-12 public school teachers. The ALP is weighted to be nationally representative on distributions of characteristics from the Current Population Survey Annual Social and Economic Supplement<sup>67</sup>. The ASLP and ATP are weighted based on distributions of characteristics from the National Center for Education Statistics' Common Core of Data<sup>68</sup>. All three panels leverage probabilistic sampling and account for non-random non-response. The study sample included 761 unique individuals from the ALP, 740 from the ASLP, and 642 from the ATP.

The Challenger Center, a non-profit organization that provides experiential STEM education programs to students throughout the United States, recruited a sample of 805 middle school students in grades six through eight. Three Challenger Learning Centers recruited students from Kentucky, Maine, and Missouri. CMU recruited 87 graduate and undergraduate students from across the university. As these two samples were convenience samples, no weights were assigned to respondents.

In total, the analytical sample contained 3,035 respondents and 12,140 respondent-video observations. Prior to data collection, the study was approved by the Institutional Review Boards of RAND, Challenger Center, and CMU and the study was performed in accordance with all relevant

guidelines and regulations. All adults provided informed consent before taking the survey and schools secured parental informed consent before fielding the survey to middle school students. Table 4 provides the descriptive statistics for each population. After the survey was complete, respondents were told of the true authenticity of the videos they viewed.

### 6.3 Identification Strategy and Analysis

Randomization of deepfake videos to respondents means that the background characteristics and beliefs of respondents were uncorrelated with the receipt of authentic or deepfake videos. Table A6 in the Appendix presents tests of baseline balance of receipt of the deepfake videos on respondent and video characteristics. Of 170 statistical tests, 15 covariates were statistically significantly imbalanced to the 10 percent level and seven to the five percent level; a rate of imbalance that is expected by chance.

We analyzed the data at the respondent-video level to obtain an overall estimate of the vulnerability of populations to STEM deepfakes and answer research question 1. We leveraged models of the following form separately on each population:

$$Y_{isr} = \beta_0 + \beta_1 \text{Deepfake}_{isr} + \alpha_i + \gamma_s + \tau_r + \varepsilon_{isr} \quad (1)$$

Where  $Y_{isr}$  was respondent,  $i$ 's, response to the video of speaker,  $s$ , randomized to be shown in the  $r$ th position (first video shown, second video shown, etc.). We analyzed three outcomes of interest: correctly identifying the authenticity of the video, incorrectly doing so, and responding "Cannot Tell." The respondent was coded as correct if they selected "Definitely Fake" or "Probably Fake" when receiving a deepfake video, or "Definitely Real" or "Probably Real" when receiving an authentic video.  $\text{Deepfake}_{isr}$  was an indicator for whether the video is a deepfake video, while  $\alpha_i$ ,  $\gamma_s$ , and  $\tau_r$  were respondent, speaker, and order fixed effects, respectively. No fixed effects were needed for identification, but were included to increase the precision of estimates. Finally,  $\varepsilon_{isr}$  was a stochastic individual-level error term. Regressions using the ALP, ASLP, and ATP were weighted to retain national representativeness and all models used clustered standard errors at the respondent level to account for the correlation between each respondent's response to the videos. We used linear probability models to ease the interpretation of the results and analyzed each population separately.

Recall that after each video we asked the respondents to indicate which of six video aspects helped them decide the video's authenticity. To answer research question 2, we first investigated whether viewing a deepfake caused respondents to differentially use aspects of the video in their decision-making process. We leveraged models of the following form:

$$A_{isr} = \beta_0 + \beta_1 \text{Deepfake}_{isr} + \alpha_i + \gamma_s + \tau_r + \varepsilon_{isr} \quad (2)$$

Equation 2 is identical to Eq. 1 except  $A_{isr}$  was now one of the six video aspects. The coefficient of interest is  $\beta_1$ , which indicated the extent to which respondents noted a particular video aspect was used to make their decision when receiving a deepfake. If  $\beta_1$  is small in magnitude and statistically insignificant, respondents are equally likely to have considered that video aspect in their decision-making progress, regardless of the authenticity of the video they are viewing. If  $\beta_1$  is positive (negative) and statistically significant, respondents are more (less) likely to have considered that aspect of a video when receiving a deepfake video.

We then investigate whether considering a video aspect leads to a higher likelihood of correctly detecting the video's authenticity. We leveraged models of the following form:

$$Y_{isr} = \beta_0 + \mathbf{A}_{isr} \mathbf{\beta}_1 + \alpha_i + \gamma_s + \tau_r + \varepsilon_{isr} \quad (3)$$

Equation 3 is the same as Eq. 1 except the indicator for obtaining a deepfake was replaced with a vector of the six video aspects,  $\mathbf{A}_{isr}$ . We estimated results separately for the subsamples of authentic and deepfake videos. The vector of coefficients of interest is  $\mathbf{\beta}_1$ , which estimated the relationship between analyzing each video aspect and the probability of correctly, incorrectly, or being unable to identify the video's authenticity. Because video aspects were not randomly assigned, these results were correlational.

Finally, we explore which subsets of populations are more vulnerable to deepfake videos to answer research question 3. We leveraged models of the following form:

$$Y_{isr} = \beta_0 + \beta_1 \text{Deepfake}_{isr} + \beta_2 \text{Deepfake}_{isr} * X_i + \alpha_i + \gamma_s + \tau_r + \varepsilon_{isr} \quad (4)$$

Equation 4 was identical to Eq. 1 except we included an interaction term of the indicator for viewing the deepfake with the respondent characteristic of interest,  $X_i$ . The coefficient of interest is  $\beta_2$ , which provides the differential effect of receiving the deepfake video for the subgroup of interest compared to the reference group. We test the constructs previously detailed in the *Survey Instrument* section and the effect of the order of seeing the video. The larger number of individual and video characteristics could have presented a multiple hypothesis comparison problem. We did not formally correct for multiple hypotheses due to the exploratory nature of the analyses, but we guarded against it by only presenting results that showed a pattern of statistical significance across populations. In addition, we explored results from models that removed respondent fixed effects and include a vector of

covariates of the characteristics in Table 4. We present results from both models to show stability of point estimates and robustness of statistical significance.

**Table 4:** Descriptive Statistics

	Adults			Principals			Teachers			Middle School Students			CMU Students		
	Mean	St. Err.	N	Mean	St. Err.	N	Mean	St. Err.	N	Mean	St. Err.	N	Mean	St. Err.	N
Male	0.481	0.029	761	0.486	0.019	735	0.234	0.017	641	0.447	0.018	805	0.519	0.057	79
Age	52.614	0.898	761	49.179	0.299	739	44.172	0.425	642	12.653	0.038	792	20.291	0.17	79
White	0.756	0.026	761	0.643	0.018	740	0.684	0.019	642	0.862	0.012	805	0.354	0.054	79
Black	0.117	0.02	761	0.092	0.011	740	0.065	0.01	642	0.073	0.009	805	0.038	0.022	79
Hispanic	0.165	0.021	761	0.07	0.01	740	0.23	0.017	642	0.042	0.007	805	0.089	0.032	79
Other	0.127	0.021	761	0.098	0.011	740	0.09	0.012	642	0.065	0.009	805	0.608	0.055	79
Did Not State Race	0	0	761	0.167	0.014	740	0.161	0.015	642	0	0	805	0	0	79
HS or less	0.365	0.032	761	--	--	--	--	--	--	--	--	--	0.089	0.032	79
Some college	0.138	0.014	761	--	--	--	--	--	--	--	--	--	0.608	0.055	79
Employed	0.638	0.026	761	--	--	--	--	--	--	--	--	--	--	--	--
Married	0.644	0.028	761	--	--	--	--	--	--	--	--	--	--	--	--
Science Occupation	0.123	0.016	761	--	--	--	--	--	--	--	--	--	--	--	--
Bachelor's Degree	0.121	0.016	761	0.01	0.004	740	0.347	0.02	642	--	--	--	0.304	0.052	79
Masters or More	0.174	0.017	761	0.988	0.004	740	0.653	0.02	642	--	--	--	--	--	--
Born in US	0.884	0.019	761	--	--	--	--	--	--	--	--	--	0.722	0.051	79
<b>Educator Variables</b>															
Years of Experience	--	--	--	11.224	0.268	737	16.182	0.348	642	--	--	--	--	--	--
Title I School	--	--	--	4.651	0.056	708	4.888	0.052	632	--	--	--	--	--	--
Elementary School	--	--	--	0.487	0.019	711	0.477	0.021	633	--	--	--	--	--	--
Middle School	--	--	--	0.229	0.016	711	0.207	0.016	633	--	--	--	--	--	--
High School	--	--	--	0.285	0.018	711	0.315	0.018	633	--	--	--	--	--	--
Math/Science Teacher	--	--	--	--	--	--	0.193	0.016	642	--	--	--	--	--	--
Social Science Teacher	--	--	--	--	--	--	0.068	0.01	642	--	--	--	--	--	--
Other Teacher	--	--	--	--	--	--	0.739	0.017	642	--	--	--	--	--	--

Notes: Statistics from the adult, principal, and teacher populations are weighted to be nationally representative.

## Declarations

### Acknowledgements and Funding Sources

The authors would like to thank the RAND Corporation American Life Panel/American Educator Panels staff for their support in developing and fielding surveys. This work was funded by the National Science Foundation (Grant #s 2039612, 2039613, and 2039614).

## Author Contributions

JM, CD, and DS led the design of the survey instruments; CD and TW analyzed data; DS and CT created the deepfakes; CD, JM, DS, DK, VAFK, LB and CT conducted data collection; JM and CD wrote the paper with input from all authors.

## Data Availability Statement

The data are not yet publicly available due to an embargo but can be made available to Scientific Reports upon request to aid in the review of this submission. Please email Christopher Doss at [cdoss@rand.org](mailto:cdoss@rand.org) to request the data. A deidentified version of the data will be publicly available in the summer of 2022.

## Additional Information

The authors declare no competing interests.

## References

1. D. Acemoglu, D. Laibson, J. A. List, Equalizing Superstars: The Internet and the Democratization of Education. *American Economic Review* **104**, 523-527 (2014).
2. O. B. Adedoyin, E. Soykan, Covid-19 pandemic and online learning: the challenges and opportunities. *Interactive Learning Environments*, 1-13 (2020).
3. E. Tan, Informal learning on YouTube: exploring digital literacy in independent online learning. *Learning Media and Technology* **38**, 463-477 (2013).
4. A. Gutierrez-Martin, A. Torrego-Gonzalez, M. Vicente-Marino, Media education with the monetization of YouTube: the loss of truth as an exchange value. *Cultura Y Educacion* **31**, 267-295 (2019).
5. J. S. Radesky, A. Schaller, S. L. Yeo, H. M. Weeks, M. B. Robb (2020) Young kids and YouTube: How ads, toys, and games dominate viewing. (Common Sense Media, San Francisco, CA).
6. T. Kind, Y. Evans, Social media for lifelong learning. *Int Rev Psychiatry* **27**, 124-132 (2015).
7. M. Cinelli *et al.*, The COVID-19 social media infodemic. *Scientific Reports* **10**, 10 (2020).
8. Y. X. Wang, M. McKee, A. Torbica, D. Stuckler, Systematic Literature Review on the Spread of Health-related Misinformation on Social Media. *Social Science & Medicine* **240**, 12 (2019).
9. C. Wittenberg, B. Tappin, A. J. Berinsky, D. G. Rand, The (minimal) persuasive advantage of political video over text. *Proceedings of the National Academy of Sciences of the United States of America* **118**, 7 (2021).
10. A. Goreis, M. Voracek, A Systematic Review and Meta-Analysis of Psychological Research on Conspiracy Beliefs: Field Characteristics, Measurement Instruments, and Associations With Personality Traits. *Frontiers in Psychology* **10**, 13 (2019).
11. S. M. Sylvester, COVID-19 and Motivated Reasoning: The Influence of Knowledge on COVID-Related Policy and Health Behavior. *Social Science Quarterly*, 19.
12. D. Etkin, E. Ho, Climate change: Perceptions and discourses of risk. *Journal of Risk Research* **10**, 623-641 (2007).
13. S. Loomba *et al.* (Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature Human Behavior* **5**, 337-348 (2021).
14. R.P. Curiel and H. Gonzalez-Ramirez. Vaccination strategies against COVID-19 and the diffusion of anti-vaccination views. *Scientific Reports*, **11** (2021).
15. M. Del Vicario *et al.*, The spreading of misinformation online. *Proceedings of the National Academy of Sciences of the United States of America* **113**, 554-559 (2016).
16. A. Photiou, C. Nicolaides, and P.S. Dhillon Social status and novelty drove the spread of online information during the early stages of COVID-19. *Scientific Reports* **11** (2021).
17. D. A. Scheufele, N. M. Krause, Science audiences, misinformation, and fake news. *Proc Natl Acad Sci U S A* **116**, 7662-7669 (2019).
18. V. F. Reyna, A scientific theory of gist communication and misinformation resistance, with implications for health, education, and policy. *Proceedings of the National Academy of Sciences of the United States of America* **118**, 6 (2021).
19. S. Tulyakov, M. Y. Liu, X. D. Yang, J. Kautz, Ieee (2018) MoCoGAN: Decomposing Motion and Content for Video Generation. in *31st IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (Ieee, Salt Lake City, UT), pp 1526-1535.
20. Y. Mirsky, W. Lee, The Creation and Detection of Deepfakes: A Survey. *Acm Computing Surveys* **54**, 41 (2021).
21. S. Lewandowsky, U. K. H. Ecker, J. Cook, Beyond Misinformation: Understanding and Coping with the "Post-Truth" Era. *Journal of Applied Research in Memory and Cognition* **6**, 353-369 (2017).

22. A. Y. Ng, M. I. Jordan, On Discriminative vs. Generative classifiers: A comparison of logistic regression and naive Bayes. *Advances in Neural Information Processing Systems 14, Vols 1 and 2* **14**, 841-848 (2002).
23. C. Zhao *et al.*, Dermoscopy Image Classification Based on StyleGAN and DenseNet201. *Ieee Access* **9**, 8659-8679 (2021).
24. Z. H. Wu, G. S. Lin, Q. Y. Tao, J. F. Cai, Acm, M2E-Try On Net: Fashion from Model to Everyone. *Proceedings of the 27th Acm International Conference on Multimedia (Mm'19)*, 293-301 (2019).
25. A. Siarohin, S. Lathuiliere, S. Tulyakov, E. Ricci, N. Sebe, First Order Motion Model for Image Animation. *Advances in Neural Information Processing Systems 32 (Nips 2019)* **32**, 11 (2019).
26. S. R. A. Ahmed, E. Sonuc, Deepfake detection using rationale-augmented convolutional neural network. *Applied Nanoscience*, 9.
27. A. Chinthath *et al.*, Recurrent Convolutional Structures for Audio Spoof and Video Deepfake Detection. *Ieee Journal of Selected Topics in Signal Processing* **14**, 1024-1037 (2020).
28. T. Zhao *et al.*, Learning self-consistency for deepfake detection. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15023-15033 (2021).
29. L. Guarnera, O. Giudice, S. Battiato, Deepfake detection by analyzing convolutional traces. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 666-667 (2020).
30. Y. Nirkin, L. Wolf, Y. Keller, T. Hassner, DeepFake Detection Based on Discrepancies Between Faces and their Context. *IEEE Trans Pattern Anal Mach Intell* **PP** (2021).
31. H. Zhao *et al.*, Multi-attentional deepfake detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2185-2194 (2021).
32. T. Mittal, U. Bhattacharya, R. Chandra, A. Bera, D. Manocha (2020) Emotions Don't Lie. in *Proceedings of the 28th ACM International Conference on Multimedia*, pp 2823-2832.
33. K. Chugh, P. Gupta, A. Dhall, R. Subramanian (2020) Not made for each other- Audio-Visual Dissonance-based Deepfake Detection and Localization. in *Proceedings of the 28th ACM International Conference on Multimedia*, pp 439-447.
34. D. Cozzolino, A. Rössler, J. Thies, M. Nießner, L. Verdoliva, Id-reveal: Identity-aware deepfake video detection. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15108-15117 (2021).
35. S. Das, S. Seferbekov, A. Datta, M. S. Islam, M. R. Amin, Towards Solving the DeepFake Problem: An Analysis on Improving DeepFake Detection using Dynamic Face Augmentation. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3776-3785 (2021).
36. X. Li *et al.* (2020) Sharp Multiple Instance Learning for DeepFake Video Detection. in *Proceedings of the 28th ACM International Conference on Multimedia*, pp 1864-1872.
37. L. Trinh, M. Tsang, S. Rambhatla, D. Y. Liu, Interpretable and trustworthy deepfake detection via dynamic prototypes. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1973-1983 (2021).
38. P. P. Yu, Z. H. Xia, J. W. Fei, Y. J. Lu, A Survey on Deepfake Video Detection. *Iet Biometrics* **10**, 607-624 (2021).
39. D. J. Robertson *et al.*, Detecting morphed passport photos: a training and individual differences approach. *Cognitive Research-Principles and Implications* **3**, 11 (2018).
40. J. Choudrie *et al.*, Machine learning techniques and older adults processing of online information and misinformation: A covid 19 study. *Computers in Human Behavior* **119**, 11 (2021).
41. M. Groh, Z. Epstein, C. Firestone, R. Picard, *Proc Natl Acad Sci U S A* **119** (2022).
42. P. J. Phillips *et al.*, Face recognition accuracy of forensic examiners, superrecognizers, and face recognition algorithms. *Proceedings of the National Academy of Sciences of the United States of America* **115**, 6171-6176 (2018).
43. G. Levchuk, A. Fouse, K. Pattipati, D. Serfaty, R. McCormack (2018) Active Learning and Structure Adaptation in Teams of Heterogeneous Agents: Designing Organizations of the Future. in *Conference on Next-Generation Analyst VI* (Spie-Int Soc Optical Engineering, Orlando, FL).
44. J. T. Wixted, L. Mickes, J. C. Dunn, S. E. Clark, W. Wells, Estimating the reliability of eyewitness identifications from police lineups. *Proceedings of the National Academy of Sciences of the United States of America* **113**, 304-309 (2016).
45. D. J. Watts, D. M. Rothschild, M. Mobius, Measuring the news and its impact on democracy. *Proceedings of the National Academy of Sciences of the United States of America* **118**, 6 (2021).
46. A. S. Tseng, Students and evaluation of web-based misinformation about vaccination: critical reading or passive acceptance of claims? *International Journal of Science Education Part B-Communication and Public Engagement* **8**, 250-265 (2018).
47. N. R. Goulden, The roles of national and state standards in implementing speaking, listening, and media literacy. *Communication Education* **47**, 194-208 (1998).
48. R. Hobbs, A review of school-based initiatives in media literacy education. *American Behavioral Scientist* **48**, 42-59 (2004).
49. N. Mirra, A. Garcia, In Search of the Meaning and Purpose of 21st-Century Literacy Learning: A Critical Review of Research and Practice. *Reading Research Quarterly* **56**, 463-496 (2021).

50. A. J. Sharon, A. Baram-Tsabari, Can science literacy help individuals identify misinformation in everyday life? *Science Education* **104**, 873-894 (2020).
51. C. Vaccari, A. Chadwick, Deepfakes and Disinformation: Exploring the Impact of Synthetic Political Video on Deception, Uncertainty, and Trust in News. *Social Media + Society* **6**, 13 (2020).
52. T. Dobber, N. Metoui, D. Trilling, N. Helberger, C. de Vreese, Do (Microtargeted) Deepfakes Have Real Effects on Political Attitudes? *International Journal of Press-Politics* **26**, 69-91 (2021).
53. M. J. Hornsey *et al.*, Evidence for motivated control: Understanding the paradoxical link between threat and efficacy beliefs about climate change. *Journal of Environmental Psychology* **42**, 57-65 (2015).
54. K. M. D. Treen, H. T. P. Williams, S. J. O'Neill, Online misinformation about climate change. *Wiley Interdisciplinary Reviews-Climate Change* **11**, 20 (2020).
55. C. B. Cooper, Media Literacy as a Key Strategy toward Improving Public Acceptance of Climate Change Science. *Bioscience* **61**, 231-237 (2011).
56. B. Zhao, S. Z. Zhang, C. X. Xu, Y. F. Sun, C. B. Deng, Deep fake geography? When geospatial data encounter Artificial Intelligence. *Cartography and Geographic Information Science* **48**, 338-352 (2021).
57. C. H. Shen *et al.*, Fake images: The effects of source, intermediary, and digital media literacy on contextual assessment of image credibility online. *New Media & Society* **21**, 438-463 (2019).
58. F. Azevedo, J. T. Jost, The ideological basis of antiscientific attitudes: Effects of authoritarianism, conservatism, religiosity, social dominance, and system justification. *Group Processes & Intergroup Relations* **24**, 518-549 (2021).
59. J. Langa, Deepfakes, real consequences: Crafting legislation to combat threats posed by deepfakes. *Boston University Law Review* **101**, 761-801 (2021).
60. W. J. Pan, D. Y. Liu, J. Fang, An Examination of Factors Contributing to the Acceptance of OnlineHealth Misinformation. *Frontiers in Psychology* **12**, 11 (2021).
61. M. W. Prull, M. B. Yockelson, Adult Age-Related Differences in the Misinformation Effect for Context-Consistent and Context-Inconsistent Objects. *Applied Cognitive Psychology* **27**, 384-395 (2013).
62. J. Saunders, A. Jess, The effects of age on remembering and knowing misinformation. *Memory* **18**, 1-11 (2010).
63. S. Lewandowsky, "Climate Change Disinformation and How to Combat It" in Annual Review of Public Health, Vol 42, 2021, J. E. Fielding, Ed. (Annual Reviews, Palo Alto, 2021), vol. 42, pp. 1-21.
64. T. C. O'Brien, R. Palmer, D. Albarracin, Misplaced trust: When trust in science fosters belief in pseudoscience and the benefits of critical evaluation. *Journal of Experimental Social Psychology* **96**, 13 (2021).
65. X. Z. Xiao, P. Borah, Y. Su, The dangers of blind trust: Examining the interplay among social media news use, misinformation identification, and news trust on conspiracy beliefs. *Public Understanding of Science* **30**, 977-992 (2021).
66. T. C. Wang *et al.*, Video-to-Video Synthesis. *Advances in Neural Information Processing Systems 31 (Nips 2018)* **31**, 13 (2018).
67. M. Pollard, M. D. Baird (2017) The RAND American Life Panel: Technical Description. (RAND Corporation).
68. M. W. Robbins, D. M. Grant (2020) RAND American Educator Panels Technical Description. (RAND Corporation).

## Figures

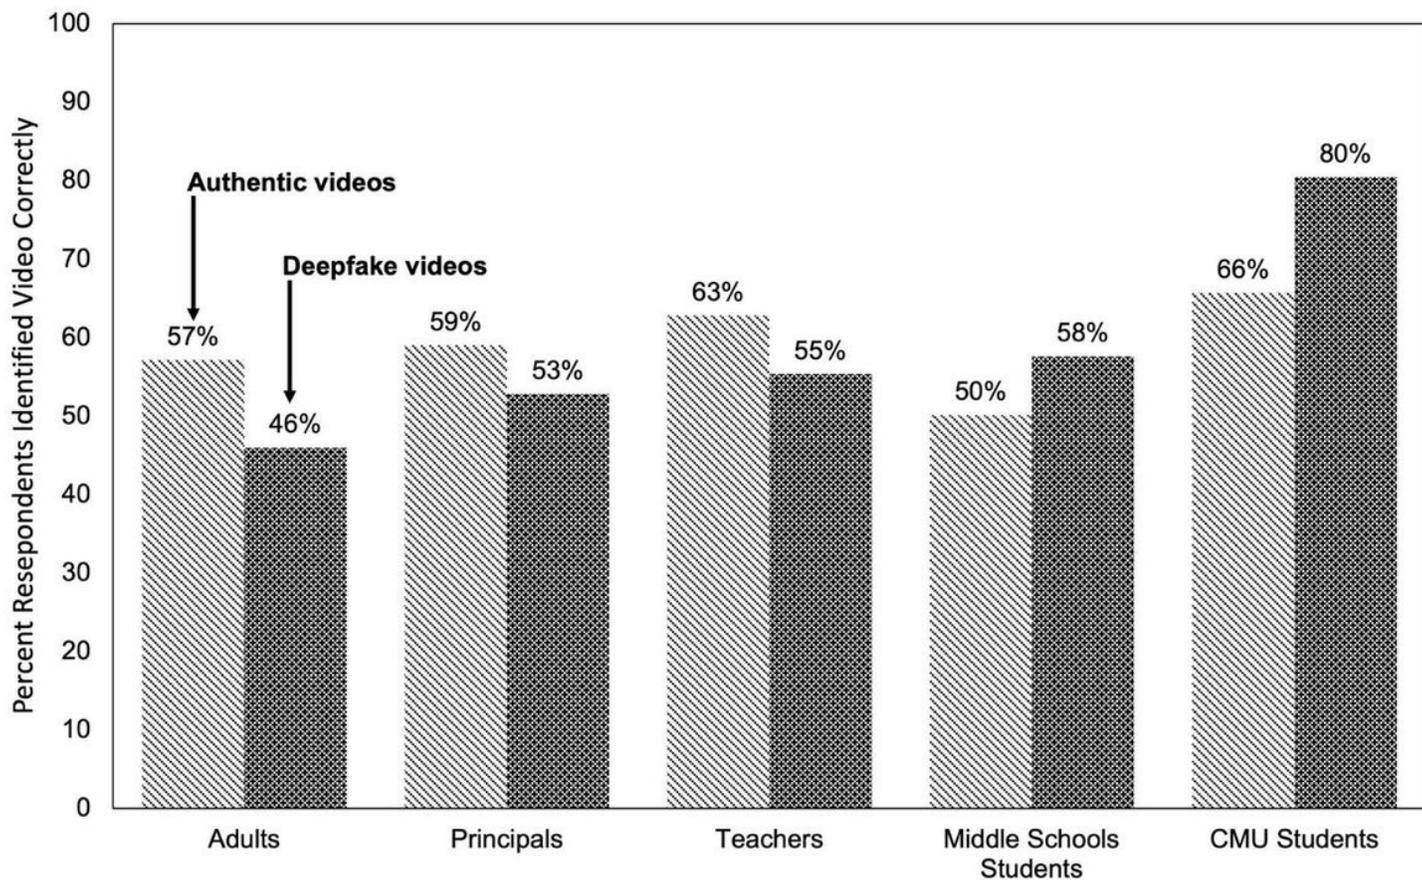


Figure 1

**Percent of Responses Correctly Identifying the Authenticity of Videos, by Video Authenticity and Population**

Notes: Each bar represents the percentage of responses that correctly identified the authenticity of videos by population and deepfake video status. Tabulations in the adult, principal, and teacher populations are weighted to be nationally representative.

Figure 1 shows the percent of responses from each population that correctly identified authentic and deepfake videos. The light gray bars show the percentage of responses that correctly identified a video as authentic when receiving an authentic video, by population. The dark gray bars show the percentage of responses that correctly identified a video as a deepfake when presented a deepfake video, by population. Tabulations in the adult, principal, and teacher populations are weighted to be nationally representative. Statistical tests of the differences between the light gray and dark gray bars within each population are presented in Table A1 in Appendix A. All differences are statistically significant to the 1 percent level or less.

**Supplementary Files**

This is a list of supplementary files associated with this preprint. Click to download.

- [DeepfakesandScienceKnowledgeDisseminationAppendicesSR.docx](#)