

# A Representative Study on Human Detection of Artificially Generated Media Across Countries

Joel Frank  
Ruhr-Universität Bochum  
joel.frank@rub.de

Franziska Herbert  
Ruhr-Universität Bochum  
franziska.herbert@rub.de

Jonas Ricker  
Ruhr-Universität Bochum  
jonas.ricker@rub.de

Lea Schönherr  
CISPA  
schoenherr@cispa.de

Thorsten Eisenhofer  
TU Berlin  
thorsten.eisenhofer@tu-berlin.de

Asja Fischer  
Ruhr-Universität Bochum  
asja.fischer@rub.de

Markus Dürmuth  
Leibniz Universität Hannover  
markus.duermuth@itsec.uni-hannover.de

Thorsten Holz  
CISPA  
holz@cispa.de

**Abstract**—AI-generated media has become a threat to our digital society as we know it. Forgeries can be created automatically and on a large scale based on publicly available technologies. Recognizing this challenge, academics and practitioners have proposed a multitude of automatic detection strategies to detect such artificial media. However, in contrast to these technological advances, the *human perception* of generated media has not been thoroughly studied yet.

In this paper, we aim to close this research gap. We conduct the first comprehensive survey on people’s ability to detect generated media, spanning three countries (USA, Germany, and China), with 3,002 participants covering audio, image, and text media. Our results indicate that state-of-the-art forgeries are almost indistinguishable from “real” media, with the majority of participants simply guessing when asked to rate them as human- or machine-generated. In addition, AI-generated media is rated as more likely to be human-generated across all media types and all countries. To further understand which factors influence people’s ability to detect AI-generated media, we include personal variables, chosen based on a literature review in the domains of deepfake and fake news research. In a regression analysis, we found that generalized trust, cognitive reflection, and self-reported familiarity with deepfakes significantly influence participants’ decisions across all media categories.

## 1. Introduction

In his 2015 book, the historian Yuval Noah Harari wrote: “In the past, censorship worked by blocking the flow of information. In the 21st century, censorship works by flooding people with irrelevant information.” [1]. While in the original context the quote referred to *fake news*, the content of the quote is more relevant than ever. Deep generative modeling has enabled the unbounded creation of fake media at scale. While it has been used for harmless endeavors like putting Jim Carrey into the movie *Shining* [2], there are also more destructive examples like phishing \$243,000 from a UK company by imitating their CEO’s voice [3] or

influencing political events [4]–[7]. A prominent example being a fake video of Ukraine’s president Zelenskyy telling his forces to surrender [8]. These fakes have resulted in a flurry of techniques to automatically detect AI-generated media [9]–[22]. However, there is little research on how convincing AI-generated media is to human observers. Prior works [23]–[28] often only consider one type of media (usually images) and often rely on small sample sizes or convenience samples.

In this work, we establish the first cross-country and cross-media baseline the detection of media generated with state-of-the-art methods. In our preregistered study, we examine three different countries (USA, Germany, and China) under three different conditions (audio, image, and text), with a total number of participation of  $n = 3,002$ . The primary goal of the surveys is to answer the following three research questions: i) Can people identify state-of-the-art generated media? ii) Which demographic factors influence the identification accuracy? iii) Which cognitive factors affect the identification accuracy?

In our survey, participants are asked to rate a set of human- and machine-generated media on how believable they are. Most importantly, we find that most AI-generated samples are already so convincing that the majority of participants cannot accurately distinguish them from human-generated content. More specifically, the average detection accuracy of participants is below 50% for images and never exceeds 60% for the other media types. Moreover, we find that participants in all countries believed that most of the samples we showed to them were human-generated, compared to the 50/50 ground truth.

To evaluate which variables might improve or worsen people’s ability to detect AI-generated media, we performed a literature review of prior work in the domains of deepfake and fake news research. This review enabled us to identify multiple personal variables that might influence the decision of participants, such as cognitive reasoning [29], media literacy [24] or political orientation [30]. We included these variables in our survey. In a regression analysis, we found that generalized trust, motivated System 2 reasoning, and self-

reported familiarity with deepfakes significantly influenced people’s decision across all media categories. Furthermore, we found several other variables that influence the decision dependent on the media type.

In summary, we make the following key contributions:

- We conduct the first preregistered cross-country and cross-media survey regarding AI-generated media with more than 3,000 participants.
- Our results indicate that AI-generated media is already so convincing that the majority of participants simply guess when asked to rate them as human- or machine-generated media.
- Using a regression analysis, we show a significant influence of generalized trust, motivated System 2 reasoning, and self-reported familiarity with deepfakes across all our conditions. Additionally, we found several condition-dependent influential factors.

All information regarding the preregistration is available at <https://osf.io/xy6v5>. The code to conduct the study and all of our analysis can be found online at [github.com/RUB-SysSec/GeneratedMediaSurvey](https://github.com/RUB-SysSec/GeneratedMediaSurvey).

## 2. Related Work

Due to the lack of comprehensive cross-media and cross-country prior work, we survey related work in the field of AI-generated media and fake news. This allows us to connect our work to prior work focusing on specific kinds of AI-generated media (e.g., images) and to comment on the transferability of methods and techniques developed in adjacent domains. Additionally, we provide an overview of the current state-of-the-art in generative modeling.

### 2.1. Generative Modeling

Generative modeling—the technology behind artificially generated media—has received tremendous attention in recent years: First, Generative Adversarial Networks (GANs) [31] started a wave of publications [32]–[42], the most prominent being several iterations of a model called StyleGAN, which, for the first time, generated photo-realistic portraits of human faces. While GANs had a profound impact, recently, the focus of the image domain shifted to diffusion models [43]–[45], the most famous being StableDiffusion [46] and Dall-E 2 [47]. These models can generate thousands of different image variants for simple text prompts with the help of large language models. Large language models are a different kind of generative models focused on generating legible text that appears to be created by humans. These models have their origin in machine translation [48]–[50], a subtask of generative modeling, where models were prompted with a paragraph in one language (e.g., English) and had to produce the equivalent in another language (e.g., French). The big breakthrough was the introduction of the attention concept [48], [51] and the corresponding architecture called Transformers [52]. Today, Transformer-based architectures like GPT-3 [53]/GPT-4 [54]

or PaLM [55] are prompted with small summaries of text and generate entire paragraphs expanding on the prompt. Finally, the synthetic speech landscape has been the most recent to be completely transformed by advances in deep learning. While traditional approaches used hand-crafted algorithms imitating human speech patterns [56]–[58], today’s algorithms use a combination of two neural networks to generate human speech from text prompts [59]–[63], so-called *Text-To-Speech* (TTS) models.

### 2.2. Personal Variables

We have organized this section according to different personal variables that have been found to influence people’s decisions.

**Media Literacy.** Previous work has suggested a link between media literacy and the susceptibility to fake information or disinformation [64]–[66]. The common assumption is that those individuals with a higher media literacy engage more critical in the media they consume. For example, a recent meta-analytic study by Jeong, Cho, and Hwang [67] indicates that people with a better understanding of media and media production systems tend to be more skeptical and realistic about media messages. Prior work has already linked the media literacy interventions to a decreased willingness to share deepfake videos [24].

**Holistic Thinking.** People from East Asian cultures tend to think more holistically, while people from Western cultures tend to think more analytically [68]–[71]. East Asians focus more on the relationship between objects and the field to which it belongs. In contrast, Westerners apply a more analytic style, focusing their attention more on an object itself [71]. In the realm of fake news, previous work has shown a negative correlation between analytic thinking and perceived accuracy of fake news [72], [73].

**Generalized Trust.** Trust is a core element of society, and everyday social life would be impossible without it [74], [75]. It has been defined as a willingness to be vulnerable to the actions of others [76] and arises from social attitudes regarding the world and other people. These attitudes can be developed either in a general or interpersonal context [75]. General or depersonalized trust refers to one’s trust towards public institutions or out-group members [77]. Past research has found connections between trust and favorable national outcomes, such as economic growth and earnings, as well as a variety of desirable interpersonal qualities (e.g., social solidarity, tolerance, volunteerism, cooperation, optimism; [78]–[80]). It is also crucial for online social interactions like online dating [81], [82] or marketplaces like Airbnb [83], [84]. In the context of machine-generated media, Mink, Luo, Barbosa, *et al.* [25] have shown that people with a higher general trust are more likely to trust and accept friend requests from machine-generated LinkedIn profiles.

**Cognitive Reflection.** The dual process theory is a cognitive framework that proposes two distinct modes of thinking: *System 1* and *System 2*. System 1 acts automatically

and spontaneously and does not require conscious reflection, while System 2 is believed to require deliberation, analytic thinking, and concentration [85], [86]. We hypothesize that people who engage more in System 2 reasoning might be better at recognizing deepfakes because they take a more conscious approach to evaluating media instead of relying on gut decisions. We administer the Cognitive Reflection Test (CRT) [86], which measures a person’s ability to engage in System 2 reasoning. Previous work found evidence that achieving a high score on the CRT is positively correlated to correctly discerning fake from real news [72], [73], [87]–[90]. Moreover, prior work showed that people are not more likely to believe fake news that are consistent with their own political ideology [87], [88], [90]. In other words, people fall for fake news because of “lazy thinking” and not because it corresponds to their beliefs [87]. A real-world analysis of Twitter data also suggests that people with higher CRT scores are more likely to share high-quality content and therefore do not support the dissemination of fake news [91]. For deepfakes, previous works have made similar findings: Cognitive reflection is positively correlated to the ability to detect artificially generated media [29], [92] and negatively correlated to inadvertent sharing of deepfakes [30], [93]. However, in a study including a political deepfake video, Appel and Prietzel [29] have shown that although a higher CRT score helps identify implausible content, it does not have an influence on the ability to spot generation artifacts (e.g., glitches or out-of-sync lips).

**Political Orientation.** Previous work suggests a link between political interest and orientation, and fake news and deepfake engagement (see [30], [94], [95]). For example, participants with higher political interest from both the US and Singapore were more likely to unintentionally share deepfakes [30]. More exposure to deepfakes was associated with being more likely to share deepfakes inadvertently, with US participants reporting more exposure to deepfakes than participants from Singapore. Chadwick and Vaccari [95] also found that people in the UK with a higher political interest were more likely to share exaggerated or false information on social media, both intentionally and unintentionally. In one of two studies with deepfake video(s) of politicians by Appel and Prietzel [29], political interest was positively associated with the likelihood of detecting the deepfake. However, in the second study this effect was not found [29]. In addition to (increased) political interest, conservatism was found to be related to fake news identification. Calvillo, Ross, Garcia, *et al.* [94] found that conservatism and news discernment, in this case headlines related to COVID-19, correlated negatively.

### 3. Method

To obtain a comprehensive overview of people’s ability to detect deepfakes, we conducted an online survey in the USA, Germany, and China between June and September 2022. These three countries were chosen for three reasons: i) availability of high-quality generative models, ii) the ability to obtain a sufficiently large qualitative sample, and iii) a

diversity of cultures. To conduct the survey, we assigned the participants to three groups and confronted them with either audio files, images, or texts. We then asked our participants to rate each stimulus as human- or machine-generated. We have preregistered the whole study design, including the sampling and analysis plan, via the Open Science Framework (OSF) before conducting the study.

#### 3.1. AI-generated Media

As the basis for our study, we collect and generate human- and machine-generated media for the three considered conditions. In the following, we describe the conducted process to obtain the audio, image, and text stimuli. Examples of all media types can be found in Appendix C.

**3.1.1. Audio.** Speech samples are generated by a text-to-speech (TTS) pipeline based on two components: the acoustic model, specifically a Tacotron 2 model [96], and the vocoder, in form of Hifi-GAN [62]. Both are state-of-the-art models used in current TTS pipelines. The first component transforms text into a Mel spectrogram representation, and the second component creates a raw waveform using the Mel spectrogram as input. For each language, specifically trained models are utilized, using the following three datasets: For English we use the LJSpeech dataset [97], for Chinese the CSMSC dataset [98], and for German the HUI dataset [99]. For the generated files, we randomly picked a disjoint set of 15 samples and generated corresponding machine-generated versions with the generative models for the respective language, using the original text as input. For the human comparison, we randomly picked another 15 samples from the corresponding training set.

**3.1.2. Image.** We use a subset of the real and machine-generated images used by Nightingale and Farid [26]. They generated synthetic faces using StyleGAN2 [42] and created a collection of 400 images that are equally split across gender, age, and ethnicity (African American or Black, Caucasian, East Asian, South Asian). Additionally, they manually checked synthetic images for uniform backgrounds and no obvious rendering artifacts. Each artificial image was matched to a real image from the training dataset of StyleGAN2, using its feature representation obtained by a deep convolutional neural network [100]. Specifically, the network extracted a 4,096-dimensional vector, which was then compared to those of all 70,000 real faces using Euclidean distance. They manually picked the most similar image from the best matches, taking position, posture, facial expression, and presence of accessories into account. For this survey, we picked 16 pairs of real and machine-generated images, ensuring an equal distribution of gender and ethnicity.

**3.1.3. Text.** We generate fake news articles with state-of-the-art language models for text generation. Real articles are collected from a major press agency in the respective country: National Public Radio (NPR) [101] in the USA, Tagesschau [102] in Germany, and China Central Television

(CCTV) [103] in China. All agencies are rated with a neutral media bias [104] (i.e., they are neither left or right oriented). To make news articles comparable across countries, we consider articles from the business, national, and international category.

For generation, we follow prior work and use a few-shot learning approach [53], in which a pre-trained model is fine-tuned for a specific task by presenting it with few examples of the desired output (i.e., a news article). We generate a sample by prompting the model with two news articles composed of a title, short summary, and the main text. The final sample is then generated by providing the model only with the title and summary of an article and outputting the main text by the model. We use the Davinci GPT3 model from OpenAI [105], which is capable of generating English, German, as well as Chinese text. We configure the model with its default parameters. To obtain a larger variety of text and avoid that the model is repeating verbatim from the summary or title, we set the presence penalty to 2, the frequency penalty to 2, and the temperature to 1. For the generated texts, we target a length of 90-100 words for German and English, respectively, and 130-140 characters for Chinese. After generation, we use a set of heuristics for syntactic post-processing (e.g., removing unnecessary white space, adding missing punctuation, or unify the formatting of time and dates). Critical, none of these heuristics change the semantics of the generated text. In total, we collected 30 news articles in each country (10 per category). For each of these articles, we computed up to five fakes as described above. From the resulting corpus of articles, we randomly selected 15 real and fake texts such that these are balanced across categories and article length.

### 3.2. Questionnaire

We implemented a custom framework for collecting the responses from participants with regard to the specific requirements for displaying text, image, and audio. Appendix B shows an example view from this survey. During the study, participants were randomly assigned to one of three media types (image, audio, and text). After answering standard demographic questions—such as age, education, and gender—participants were asked about their knowledge on *deepfakes* with a closed question (5-point Likert scale). In the next segment, we showed samples of human- and machine-generated media in a per-participant randomized order, and the participants were asked to rate how believable they were on a scale from  $-3$  (definitely non-human) to  $+3$  (definitely human). Each participant saw 50% real and 50% fake media. This ratio was not disclosed to the participants. Participants were only informed that the data contained human- and machine-generated data. We showed the participants 15 human- and 15 machine-generated audio or text samples. For images, we showed a split of 16/16 since we balanced the dataset for ethnicity and gender (cf. Section 3.1.2). After this experimental part of the questionnaire, the participants answered questions for several standard scales, which correspond to the variables of interest described in Section 2.2: the

New Media Literacy Scale (NMLS) [106] (media literacy), the cross-cultural version of the Analysis-Holism Scale (AHS) [107] (holistic thinking), the Generalized Trust Score (GTS) [108] (generalized trust), the Inglehart index [109] (political orientation), and the CRT [86] (cognitive reflection). The order of the post-experiment scales were randomized per participant. All questions were mandatory, and participants could not continue without answering all questions.

The English and German translations of the survey were done by us. The Chinese version was done by our University’s department of East Asian studies. To meet the challenges of a multilingual survey, we asked native speakers to back-translate some questions and assess the quality of the translation.

### 3.3. Data Collection and Participant Compensation

We obtained approval from our institution’s Institutional Review Board (IRB), and our study protocol was deemed to comply with all ethical regulations. In the period from June to September 2022, our online survey was distributed in the three countries (United States, Germany, and China) by the panel provider Lightspeed Research (Kantar). Kantar handled participant recruitment, country representative quotas, and participant compensation and has committed itself to follow the ICC/ESOMAR code of conduct [110]. Our panel provider did not disclose the actual participant compensation to us, and we had no influence on the compensation. They calculated a cost of 2.70€ per completed survey, which might be below the legal minimum wage in at least one of the countries surveyed. However, this compensation is in a similar range to that of crowdworking platforms [111], and is—according to our panel provider—in line with industry standards. We cannot verify this statement, as we do not have comparable data on online panelists’ compensation.

At the beginning of the study, participants were informed in detail about the purpose of the study and the use of their data. We also informed participants that they can cancel their participation in the study at any time up until the end of the survey. Afterwards, data withdrawal is no longer possible as participation is anonymous. We removed any incomplete data from the dataset. Participants gave their informed consent by checking a checkbox and clicking a button titled “submit”.

To enhance answer quality, we implemented an attention check question that asked participants to select one specific response, e.g., “This is an attention check please select: Neutral”. All participants that failed our attention check were removed. Overall, our sample consists of  $n = 3,002$  participants (USA:  $n = 1,001$ ; Germany: 1,000; China:  $n = 1,001$ ).

### 3.4. Data Cleaning

Following Meade and Craig [112], we looked at the overall time a participant took to complete the survey and discarded every participant outside the 95% interval. Note that we stratified the computed interval by country and media category, e.g., we computed different intervals for

China/audio, China/image, and USA/audio. Additionally, we discarded every participant who rated every media the same, e.g., everything with “definitely human”. Finally, our total number of eligible participants was  $n = 2,609$  (USA = 822; Germany = 875; China = 912).

## 4. Results

We follow our preregistration and perform an exploratory and regression analysis on the data. We analyze the raw ratings of the participants as well as derived accuracy scores (cf. Section 3.4). We perform the exploratory analysis for both ratings and accuracy scores but, differing from our preregistration, only run a regression analysis on the latter. Originally, we planned to also perform a full analysis of the ratings, but due to space constraints, we are leaving this as an open question for future work.

### 4.1. RQ1: Can People Identify SOTA Generated Media?

We report a summary of our dataset in Table 1. Generally speaking, participants in all countries and across all media types predominantly rated media human-generated, regardless of whether the underlying media was actually created by a human or a machine. The average rating peaks at 1.81 for audio in the US and bottoms out at 0.17 for audio in China. Kruskal-Wallis one-way ANOVA showed a statistically significant difference across all media for both human-generated (audio:  $H(2) = 339.80$   $p < 0.001$ ; image:  $H(2) = 142.60$   $p < 0.001$ ; text:  $H(2) = 87.19$   $p < 0.001$ ) and machine-generated media (audio:  $H(2) = 331.56$   $p < 0.001$ ; image:  $H(2) = 257.55$   $p < 0.001$ ; text:  $H(2) = 143.99$   $p < 0.001$ ). We perform post-hoc Mann-Whitney-U tests (Bonferroni corrected) and report the results in Table 3. We find statistically significant differences between all pairs, except for USA-Germany (Human-Audio), Germany-China (Machine-Audio), USA-Germany (Machine-Image), and USA-China (Human-Text).

In Table 2, we investigate these differences in more detail and report the percentage of human-rated samples per country and media type. In general, all media, across all categories, were predominantly rated as human-generated, even if they were artificially created. US participants show a larger tendency to predict a sample as being human-generated (roughly between 70-78% of the samples) independent of the category. German (between 58-76%) and Chinese (between 54-68%) participants exhibit a wider range of ratings but still predominantly rate the samples as human-generated. Looking at the results in more detail, machine-generated audio data is more often detected by German participants compared to participants from the US. While Chinese participants detect even more machine-generated data, they also more often interpret human-generated data as artificial, which leads to an overall smaller accuracy (see Table 1).

Interestingly, US as well as German participants have a tendency to rate machine-generated images more often as

TABLE 1: **Summary Statistics for the Dataset** We summarize the statistics of our dataset after filtering out ineligible participants (cf. Section 3.4). We report aggregated age and education statistics (OECD classification). Ratings are on a scale from -3 (definitely non-human) to +3 (definitely human) and are centered on 0 (unsure). AHS is measured on a 7-point likert scale, while GTS, Familiarity (FAM), and NMLS are on a 5-point scale. Scores in the CRT can range from 0-3, and the Inglehart index (Political Orientation (PO)) ranges from 1-4.

	USA ( $n = 822$ )		Germany ( $n = 875$ )		China ( $n = 912$ )	
	n	%	n	%	n	%
<b>Gender</b>						
Female	449	54.62%	447	51.09%	433	47.48%
Male	373	45.38%	428	48.91%	478	52.52%
<b>Age</b>						
18-34	175	21.29%	185	21.19%	379	41.60%
35-49	251	30.54%	218	24.97%	253	27.77%
50-64	182	22.14%	241	27.61%	256	28.10%
65+	214	26.03%	229	26.23%	23	2.52%
<b>Education</b>						
Low	27	3.28%	119	13.60%	36	3.95%
Medium	349	42.46%	479	54.74%	268	29.39%
High	446	54.26%	277	31.66%	608	66.67%
	<b>mean</b>	<b>std</b>	<b>mean</b>	<b>std</b>	<b>mean</b>	<b>std</b>
<b>Ratings (Human)</b>						
Audio	1.60	1.74	0.99	1.87	0.33	1.04
Image	1.65	1.72	0.97	1.84	0.42	0.83
Text	1.70	1.68	0.89	1.88	0.73	0.59
<b>Ratings (Machine)</b>						
Audio	1.81	1.75	0.92	1.97	0.17	0.29
Image	1.50	1.65	1.10	1.81	0.43	1.09
Text	1.77	1.75	0.76	1.89	0.49	0.27
<b>Accuracy</b>						
Audio	50.57%	10.38%	59.15%	13.98%	51.73%	11.65%
Image	48.58%	10.62%	45.90%	11.80%	49.93%	10.47%
Text	51.50%	9.40%	54.48%	10.30%	52.45%	10.22%
<b>Predictors</b>						
CRT	0.49	0.83	1.04	1.04	1.49	1.10
FAM	1.04	1.32	1.00	1.21	1.38	1.16
PO	1.30	1.03	1.20	1.04	1.12	0.86
AHS	4.73	0.56	4.97	0.62	4.89	0.49
GTS	3.44	0.74	3.42	0.66	3.81	0.65
NMLS CC	3.61	0.58	3.65	0.57	3.83	0.49
NMLS CP	2.69	1.06	2.17	0.96	3.57	0.70
NMLS FC	3.69	0.63	3.74	0.56	3.84	0.52
NMLS FP	3.01	1.04	2.72	0.91	3.80	0.60

TABLE 2: **Percentage of Samples Rated as Human-Rated** We report the percentage of samples rated as human-generated per country and media type, where we aggregate the human-generated ratings (+1, +2, +3).

	USA		Germany		China	
	Human	Machine	Human	Machine	Human	Machine
Audio	73.67%	72.20%	76.91%	58.93%	58.14%	54.93%
Image	74.23%	77.70%	70.86%	78.64%	61.14%	61.17%
Text	72.70%	70.28%	67.52%	59.27%	67.88%	62.75%

% Rated as Human

human-generated than real images, which does not hold for Chinese participants. In the case of text data, US participants are more often convinced that artificial texts are written by a human than Chinese participants. German participants are even slightly better in identifying machine-generated text. This might be due to lower quality of the generated text due to fewer training samples of the model.

TABLE 3: **Pair-wise Mann-Whitney U Tests Between the Country Ratings** We report the difference in mean ratings between different country pairs. \*  $p < 0.05$ ; \*\*  $p < 0.01$ ; \*\*\*  $p < 0.001$  - Bonferroni equivalent

**Human-Generated**

	Germany	China		Germany	China		Germany	China
USA	$0.03 \pm 2.26^*$	$0.61 \pm 2.41^{***}$	USA	$0.15 \pm 2.18^{***}$	$0.45 \pm 2.35^{***}$	USA	$0.27 \pm 2.17^{***}$	$0.09 \pm 2.36$
Germany	N/A	$0.62 \pm 2.31^{***}$	Germany	N/A	$0.31 \pm 2.35^{***}$	Germany	N/A	$-0.19 \pm 2.35^{***}$

**Machine-Generated**

	Germany	China		Germany	China		Germany	China
USA	$0.57 \pm 2.34^{***}$	$0.65 \pm 2.49^{***}$	USA	$0.03 \pm 2.12$	$0.56 \pm 2.35^{***}$	USA	$0.40 \pm 2.23^{***}$	$0.18 \pm 2.42^{***}$
Germany	N/A	$0.10 \pm 2.47$	Germany	N/A	$0.53 \pm 2.24^{***}$	Germany	N/A	$-0.23 \pm 2.37^{***}$

**Audio**

**Image**

**Text**

TABLE 4: **Pair-wise Mann-Whitney U Tests Between the Country Accuracies** We report the difference in accuracy between different country pairs. \*  $p < 0.05$ ; \*\*  $p < 0.01$ ; \*\*\*  $p < 0.01$  - Bonferroni equivalent

	Germany	China		Germany	China		Germany	China
USA	$-8.49 \pm 17.69^{***}$	$-0.40 \pm 15.43$	USA	$2.35 \pm 14.80^*$	$-1.37 \pm 14.83$	USA	$-2.58 \pm 13.60^{**}$	$-1.17 \pm 13.28$
Germany	N/A	$8.15 \pm 17.72^{***}$	Germany	N/A	$-4.06 \pm 16.23^{***}$	Germany	N/A	$1.49 \pm 14.25^*$

**Audio**

**Image**

**Text**




Figure 1: **Accuracy of the Participants** We plot the accuracy of our participants correctly identifying a media as fake or real. The levels indicate 10%, 50% and 90% accuracy.

We also analyze the overall accuracy of our participants in Figure 1. We observe the biggest difference between the groups in the audio condition (USA  $50.57 \pm 10.38$ ; Germany  $59.15 \pm 13.98$ ; China  $51.73 \pm 11.65$ ), where Germans perform better than US or Chinese participants. The difference is less pronounced for images, where Germany performs the worst (USA  $48.58 \pm 9.40$ ; Germany  $45.90 \pm 11.80$ ; China  $49.93 \pm 10.47$ ). The accuracy for text is most similar between the three groups (USA  $51.50 \pm 9.40$ ; Germany  $54.48 \pm 10.30$ ; China  $52.45 \pm 10.22$ ). Kruskal-Wallis one-way ANOVA showed a statistically significant ( $p < 0.001$ ) difference for the accuracy on audio and image, but not on text ( $p = 0.005$ ). We also run post-hoc paired Mann-Whitney U tests. The results are shown in Table 4. Germany differs for audio and image, while all countries are fairly close for text.

Finally, to better visualize our overall results, we also plot the *Receiver Operating Characteristic* (ROC) curve in Figure 2. This curve visualizes the *True Positive Rate* (TPR) against the *False Positive Rate* (FPR) of participant ratings. A TPR of 1.0 and FPR of 0.0 would indicate that every media is rated correctly. Again, we can observe that participants are mostly guessing with the exception of German audio data, which is a clear outlier.

We hypothesize that this might be due to a lower quality of the German audio data. From anecdotal evidence, the German audio samples sound more “robotic” and noisier than samples in other languages. Based on the data we have collected, we cannot make a full conclusion about the cause. However, note that the survey prompts are almost identical for the different media types, so errors in the experimental setup would have shown up in the other media types as well.

**Finding 1.** Across all media types and countries, we observe that artificially generated samples are almost indistinguishable from “real” media. Participants predominantly rated artificially generated media as human-generated and performed even worse than random guessing for images. Surprisingly, all countries are quite close to each other w.r.t. performance. This is surprising, as English media is often believed to be far ahead of media in other languages [113]. An exception is the German audio data, where the participants perform significantly better.

**4.2. Modelling Choices**

We now analyze the accuracy of our participants for human- and machine-generated media separately by running

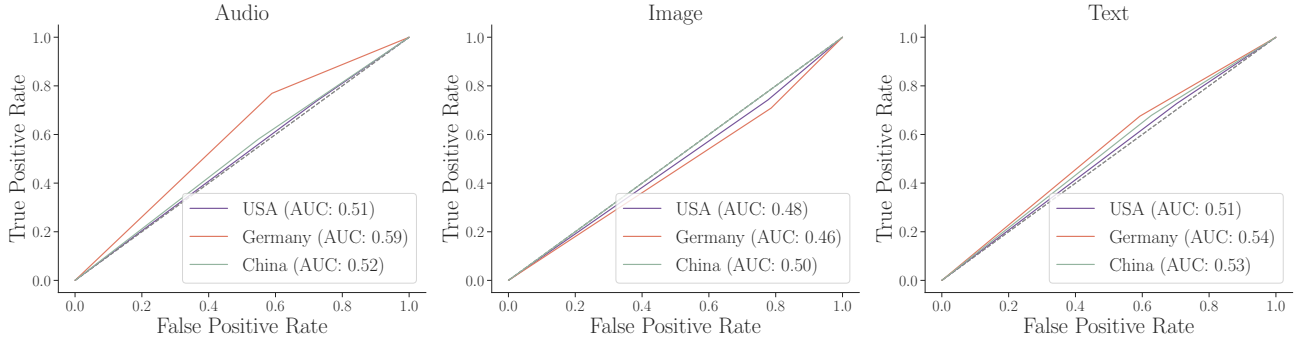


Figure 2: **ROC-Curve for Different Media Types and Country Pairs** We plot the Receiver Operating Characteristic (ROC) across media type and country pairs. We also report the Area-Under-the-Curve (AUC) values.

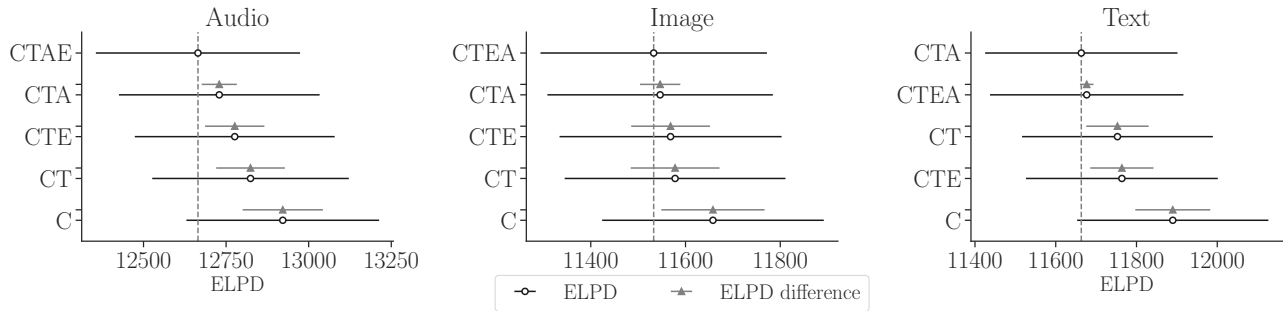


Figure 3: **Comparing the ELPD of different model iterations** We plot the ELPD of each model, as well as the expected difference for each model w.r.t. the best model. Lower ELPD values indicate a better model fit to the data. Note that a specific ELPD value is meaningless, it derives meaning from comparison with the other models' values. The figure is ordered from top to bottom, with the best performing model being on top.

The models are coded using the variables included: (C)ountry, Median (T)ime, (A)ge, (E)ducation.

a regression analysis to analyze the probability of correctly labeling a given media. Following established guidelines for analyzing cross-cultural data [114], we use a Bayesian multilevel-binomial linear mixed effect model predicting the probability of correctly classifying human- or machine-generated media. For our analysis, we choose a Bayesian framework to obtain our estimates instead of a Frequentist one. Similar models can be built with maximum likelihood estimation [115], but it has been shown that Bayesian methods produce more stable estimates [116]–[118].

Compared to the Frequentist approach, which treats model parameters as point estimates (averages), the Bayesian regression treats parameters as random variables (probability distributions). More specifically, when we use an ordinary least square regression, we want to estimate the most likely parameters; in a Bayesian regression analysis, we are instead interested in estimating a range of likely values for this parameter. This estimated distribution is referred to as the *posterior distribution*.

We choose a multilevel (or hierarchical) model to account for the different subgroups (i.e., demographic variables). Furthermore, we allow the estimates to vary at the country

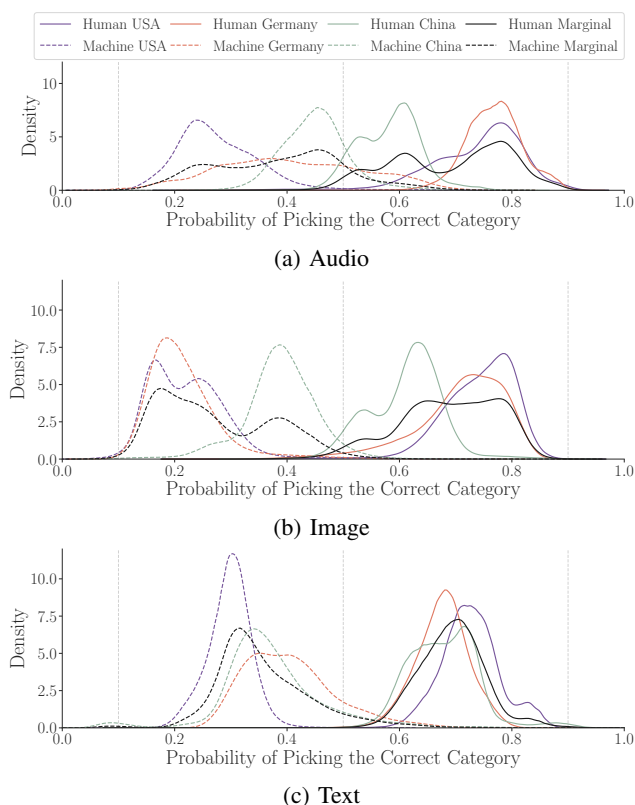
level, while also letting these results influence the population estimate. For example, the estimate for US participants in a particular age group may inform the estimate for German participants in a similar age group. An additional benefit of a Bayesian approach is that we can model locality, e.g., we can model that age groups that are closer in age are more correlated.

We first examine a comparison between different countries in Section 4.3 and choose our model to explain the differences between countries. We then build on that analysis and assess the influence of predictor values independently of the culture (i.e., on a population level) in Section 4.4.

### 4.3. RQ2: Which Demographic Factors Influence the Identification Accuracy?

We start by analyzing the influence of different demographic variables across cultures. To that end, we compare several model iterations for each media type, where each iteration adds specific demographic variables. The full model includes varying intercepts per country and three correction terms: i) one modelling the age range of a participant, ii) one





**Figure 4: Predicted Posterior Probability of the Regression Models** We plot the predicted posterior probability of picking the correct results by media type. We plot the results separated by country as well as the marginal over all countries. The levels indicate 0.1, 0.5, and 0.9.

modelling their education level, and iii) a term for the median time taken per stimulus. We partially pool these variables to obtain more stable estimates. A full model description and further details can be found in Appendix A.

To compare our models, we use the *Expected Log Pointwise Predictive Density* (ELPD), which is a measure on how well a given model generalizes to new data. We compute this metric by leave-one-out cross validation; that is, we fit a model on all but one data point, compute the log predictive densities for the left out point (i.e., measure how well the model predicts the left out point), and repeat that process for each data point. In practice, we do not refit the model each time, instead there are reliable ways to estimate this process [119].

We plot the results in Figure 3, where a smaller ELPD value indicates better model fit to the data. Additionally, we also report the expected pointwise difference for each model w.r.t. the best model found. The smaller the expected difference between the models’ ELPD, the more similar the predictions. Note that we also tried model variants omitting the time correction term, but they performed worse compared to the rest. For brevity, we omit them from the plot.

Audio and image data are best explained by the full model (CTAE). The difference is more pronounced for audio data, where including an education term is a clear advantage

over the model that only includes the age term. We analyze these differences more closely in Section 4.3.3. However, while the contrast is more stark for audio, the difference between the image models is closer. Finally, when we look at text data, the best performing model only considered median time and age of a participant.

**4.3.1. Overall Results.** In Figure 4, we plot the predicted posterior distribution of the probability of picking the correct category, both as the marginal distribution (integrated over all data) and split into the different conditional distributions given the respective countries. Note that we discuss the posterior probability predicted by our model for our sample, not the exact probabilities implied by our sample. The difference might be subtle but is important. When we discuss the exact statistics of our sample, we would assume that we have accurately sampled the underlying population, which is rarely true in reality (cf. Section 6). Instead we build a statistical model accounting for the uncertainty implied by our sampling method, which in turn allows us to better approximate the (real) underlying population [114], [119], [120].

These results give further insights into our initial rating observations presented in Section 4.1. We can again observe the trend that people are fairly good at identifying real media and are worse when it comes to identifying fakes. For example, machine-generated audio data is easier to identify (16.83% of the marginal probability mass above .5) than both image (1.28%) and text data (7.83%). This trend even persists when we discount German audio data (10.56%). The distributions also give a better perception of the variability of the participants. For example, the density for machine-generated German audio data is fairly spread out, encompassing individuals who are really good (27.39% of the probability mass above .5) and really bad (13% lower than .25) at identifying generated data.

**Finding 2.** In our preregistration, we hypothesized that image data would be the most convincing, audio to be in the middle, and text to be the least convincing. In contrast, our results suggest that fake audio data is the least convincing. We conjecture two plausible explanations: First, as introduced in Section 3.1, we generate audio data by using a text-to-speech pipeline, i.e., we supply a textual representation of the audio we want to generate. For text data, we supply the headline together with a short summary, which the model expands to a full paragraph. In hindsight, the text model has many different but equally valid ways to generate the paragraph. The audio model, on the other hand, has to both match the exact text and generate believable human voices with matching pitch and rhythm. Second, we employ the most advanced version of OpenAI’s GPT-3 models, which dwarf our speech model in both size and the volume of training data [53], [62].

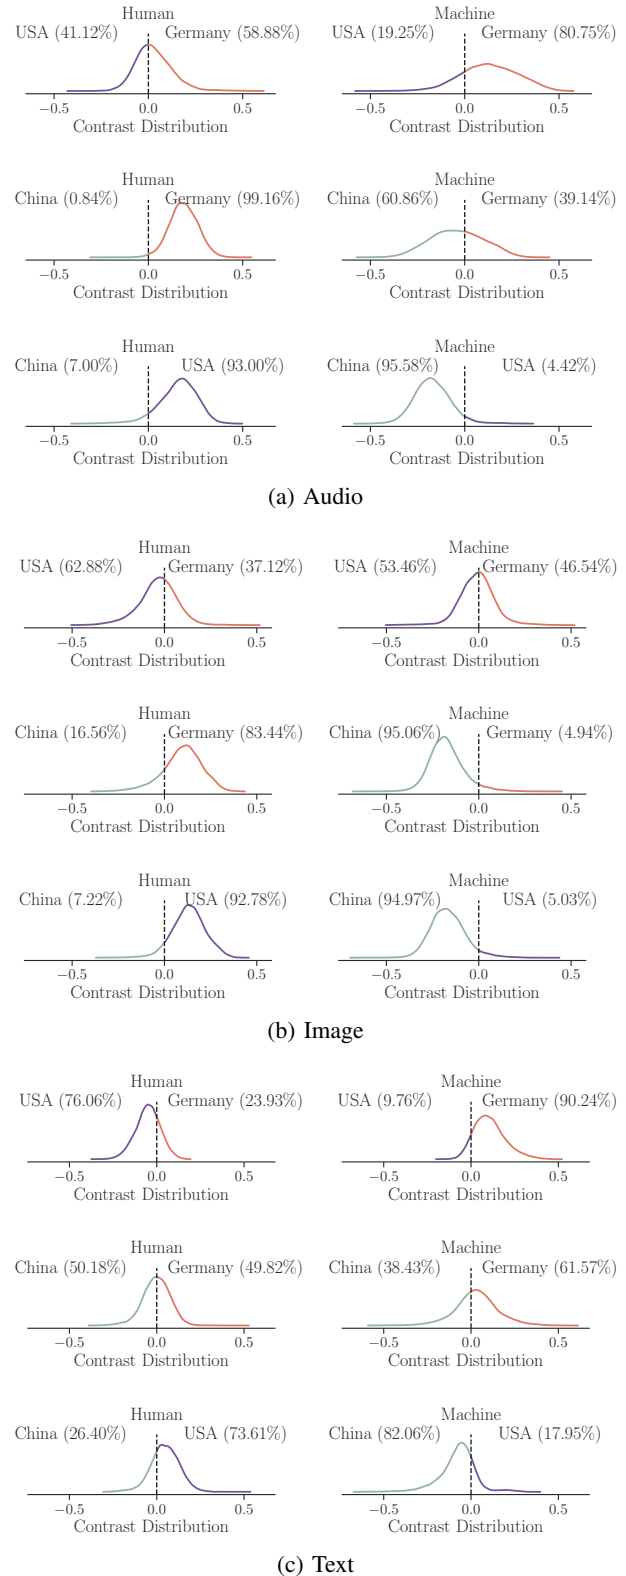


**4.3.2. Paired Comparisons.** We also performed a paired comparison between countries. To this end, we compute the contrast distributions between different country pairs, as looking at the posterior distributions directly can be deceiving [120]. To compute the contrast distributions, we randomly draw 20,000 pairs of samples from the distributions we want to compare and subtract those samples from each other to obtain a “distribution of differences”.

For example, if we want to compare participants from China and Germany for artificially generated audio, we would subtract each sample from the Chinese posterior distribution from a sample from the German posterior distribution. If German participants are generally better at identifying machine-generated media, the majority of the distribution would consist of positive values. The resulting contrast distributions are shown in Figure 5. Indeed, the contrast distribution between Germany and China for machine-generated audio samples is slightly skewed towards Chinese participants (cf. Figure 5a). That indicates that Chinese participants might have a small edge over German participants. This is not immediately obvious from Figure 4a, since German participants exhibit a fairly high variance, resulting in both people who are very good and people who are very bad at detecting machine-generated samples.

**Audio:** We start by examining audio data by inspecting the respective contrast distributions in Figure 5a. First we compare participants from Germany against participants from the USA: German participants are slightly better compared to those from the USA when detecting human-generated audio data (for 60.95% of the sampled pairs the probability of picking the correct label was higher for the German samples than for the US samples) and are clearly better at recognizing machine-generated audio samples (79.80%). As stated before, this may be explained by the lower quality of German machine-generated audio. When compared to Chinese participants, the German participants are clearly better at recognizing human-generated data (99.01%) but perform slightly worse on machine-generated audio samples (39.34%). Finally, we compare participants from the USA and China, showing that US participants are generally better at identifying human-generated samples and worse at identifying machine-generated ones. These results agree with our initial observations about the ratings: Chinese participants show a lower probability of rating a sample as human-generated, no matter if it stems from a human or a machine. While US and Chinese participants effectively randomly guess their choices, they do it at different rates, reflected by the fact that they outperform the other in different categories (while being equal overall). German participants on the other hand clearly outperform the other regions in one category while being about equal in the other.

**Image:** When comparing images, US participants perform slightly better compared to Germans on both human- (62.23% favors USA) and machine-generated samples (54.53%). The difference is bigger when comparing Germany and China, where Chinese participants are clearly better in recognizing machine-generated content (94.93% favors China), the ratio is closer for human-generated content (only



**Figure 5: Contrast Distributions Comparing Different Countries** We plot the contrast distribution obtained by subtracting one country’s posterior distribution from the other.

83.15% favors Germany). Finally, we can again find a pattern where the US and Chinese participants are distinctively better at either category. These results mirror our initial findings, where Chinese and German participants (Germans being worse) clearly differ, but the rest of the pairs are more closer together.


**Text:** For text data, we can confirm our previous observation that it yields the most similar results (cf. Section 4.1), with no clear standout. Germany and the US are the most different, with German participants outperforming the US participants more clearly on machine-generated media (86.78% of the distribution favors Germany) than the US participants do on human-generated samples (74.17%). The difference between Chinese and German participants is much smaller (59.20% and 49.80%, respectively), while the contrast distributions for the USA and China are about equal but slightly favor Chinese participants.

**Finding 3.** Our paired comparisons explain the phenomenon observed in Section 4.1 where countries often perform similarly overall, even though they perform significantly differently when compared on a per-category basis (human-/machine-generated). Examining the contrast distributions, we observe that they often seem like mirror images of each other. In other words, while one country performs better on machine-generated samples, the paired country performs better on human-generated samples. This, in-turn, leads to a similar overall accuracy. This is most noticeable for image and text and, when comparing USA and China, for audio as well.


**4.3.3. Demographic Variables.** When investigating the influence of demographic variables, we find that they only have marginal influence on Chinese and US participants. The biggest influence can be observed for German participants, who were presented with audio data. This indicates that when participants only guess, the influence of any one demographic variable tends to disappear.

We plot the influence of age split across media types in Figure 6. For readability, we plot the overall predicted posterior probability of picking the correct result. We compute this probability by keeping all other influences at average and then sampling the Gaussian Process we use to model the influence of age. The biggest impact can be spotted for German participants presented with audio files, where we can observe a steady decrease towards guessing with progressing age (mean predicted probability decreasing from 68.58% to 50.99%). Other less prominent trends are noticeable for German participants in the image condition (a decrease from 47.23% to 44.15%) and for Chinese participants in the text condition (53.28% to 50.50%).


In Figure 7, we plot the influence of the education parameter. We observe a slight increase in the mean probability for German participants in the audio condition. However, all of the posterior distributions still overlap. We also observe that



(a) Audio

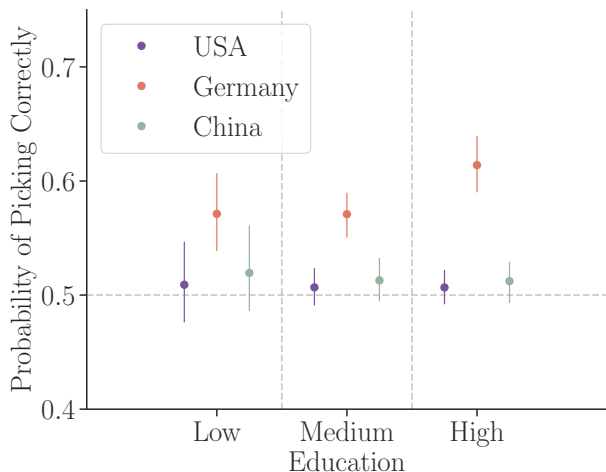


(b) Image




(c) Text

**Figure 6: Predicted Posterior Probabilities split by Age Bracket** We plot the predicted posterior probability of picking the correct result by media type. We plot the probabilities separated into the different age brackets. Note that we only show the probability range of 0.4 to 0.75 to better highlight the results. 0.5 corresponds to random guessing, higher values indicate a higher accuracy in picking the correct result.



(a) Audio



(b) Image

**Figure 7: Predicted Posterior Probabilities split by Education Levels** We plot the predicted posterior probability of picking the correct results by media type. We plot the results separated into the different education groups. Note that we only show the probability range of 0.4 to 0.75 to better highlight the results. 0.5 corresponds to random guessing, higher values indicate a higher accuracy in picking the correct result.

the distributions reflecting lower education levels exhibit a higher variability, reflecting the general trend of online panels to not accurately sample these participants (cf. Section 5).

**Finding 4.** Overall, we find that demographic variables matter—until they do not. When machine-generated media is still lagging behind in terms of quality, demographic variables can influence the rate of detection. In our case, we observe an influence of the age of participants on the detection accuracy for audio files. However, once the quality improves, the influence of the demographic variables collapses, and the probability of picking the correct result becomes almost the same across the different age/education groups.

#### 4.4. RQ3: Which Cognitive Factors Influence the Identification Accuracy?

Finally, we study the influence of different variables on the probability of classifying a given media correctly. We again perform a regression analysis with the models selected in Section 4.3 and individually add our predictor variables presented in Section 2.2. The results are presented in Table 5. We present the posterior means and 89% Prediction Interval (PI). When the majority of the posterior distribution either falls below or above zero, the predictor is believed to be meaningful. Negative scores indicate that the predictor decreases the probability of picking correct for the given category.

Note that these are prediction intervals and not confidence intervals, as Bayesian inference does not infer point estimates but ranges of plausible values. This aspect, in conjunction with hierarchical modelling, allows Bayesian methods to circumvent the problem of multiple comparisons found in Frequentist methods. Thus, we do not correct for multiple comparisons, which also aligns with remarks in Gelman, Hill, and Yajima [121].

We decided to model the parameters individually for the following reasons: Due to the link function, all terms in a generalized linear model will interact. This is true whether the estimates are obtained by Frequentist or Bayesian methods. Consequently, we cannot simply include all measured terms in a regression without running the risk of Simpson’s Paradox (cf. Pearl [122]). Intuitively, including terms in a regression can both reveal and hide causal relationships in your data. Thus, whether to include a term in the regression analysis cannot be derived from the data but needs to be based on our understanding of how the data was created. Therefore, we choose to model each predictor separately.

We found three predictor variables that were robust across all media types: First, participants with a higher general trust (GTS) have a higher chance of detecting human-generated content (audio:  $P(\beta > 0) = 1.00$ ; image:  $P(\beta > 0) = 1.00$ ; text:  $P(\beta > 0) = 0.99$ ). On the flip side, it negatively affects their ability to classify machine-generated samples (audio:  $P(\beta < 0) = 1.00$ ; image:  $P(\beta < 0) = 1.00$ ; text:  $P(\beta < 0) = 1.00$ ). Second, participants who achieved a higher CRT score are worse at recognizing human-generated media (audio:  $P(\beta < 0) = 0.98$ ; image:  $P(\beta < 0) = 1.00$ ;

TABLE 5: **Predictor Posterior Means and 89 % Prediction Interval** We display the posterior means and 89 % PI for different predictors per media categories. We highlight predictors where more than 95% of the posterior distribution falls either above or below zero in **bold**.

	Audio		Image		Text	
	Human	Machine	Human	Machine	Human	Machine
AHS	-0.04 [-0.07, -0.00]	<b>0.05 [ 0.02, 0.09]</b>	<b>-0.05 [-0.09, -0.01]</b>	0.00 [-0.04, 0.04]	<b>0.04 [ 0.00, 0.07]</b>	<b>-0.04 [-0.08, -0.01]</b>
CRT	<b>-0.14 [-0.24, -0.04]</b>	<b>0.17 [ 0.07, 0.27]</b>	<b>-0.22 [-0.31, -0.11]</b>	<b>0.17 [ 0.08, 0.28]</b>	<b>-0.09 [-0.19, -0.01]</b>	<b>0.16 [ 0.07, 0.25]</b>
FAM	<b>0.85 [ 0.48, 1.53]</b>	<b>-0.68 [-1.32, -0.36]</b>	<b>-0.36 [-0.47, -0.27]</b>	<b>0.22 [ 0.11, 0.32]</b>	<b>-0.30 [-0.38, -0.22]</b>	<b>0.18 [ 0.10, 0.26]</b>
GTS	<b>0.16 [ 0.12, 0.20]</b>	<b>-0.11 [-0.15, -0.07]</b>	<b>0.11 [ 0.07, 0.14]</b>	<b>-0.11 [-0.14, -0.07]</b>	<b>0.09 [ 0.05, 0.12]</b>	<b>-0.10 [-0.13, -0.07]</b>
NMLS CC	0.05 [-0.00, 0.09]	0.02 [-0.02, 0.07]	<b>-0.14 [-0.20, -0.09]</b>	0.04 [-0.01, 0.10]	<b>0.10 [ 0.05, 0.14]</b>	<b>-0.11 [-0.16, -0.07]</b>
NMLS CP	<b>0.09 [ 0.03, 0.14]</b>	<b>-0.19 [-0.25, -0.14]</b>	<b>0.08 [ 0.01, 0.14]</b>	<b>-0.08 [-0.15, -0.01]</b>	<b>0.13 [ 0.07, 0.19]</b>	<b>-0.10 [-0.15, -0.04]</b>
NMLS FC	<b>-0.07 [-0.11, -0.02]</b>	0.03 [-0.01, 0.08]	<b>0.22 [ 0.17, 0.27]</b>	<b>-0.23 [-0.28, -0.17]</b>	-0.03 [-0.08, 0.01]	0.03 [-0.01, 0.08]
NMLS FP	-0.02 [-0.08, 0.03]	0.02 [-0.03, 0.08]	-0.04 [-0.11, 0.02]	<b>0.07 [ 0.00, 0.14]</b>	-0.02 [-0.07, 0.04]	-0.02 [-0.08, 0.03]
PO	0.02 [-0.11, 0.14]	<b>-0.11 [-0.22, 0.00]</b>	<b>-0.41 [-0.52, -0.30]</b>	<b>0.26 [ 0.15, 0.38]</b>	0.08 [-0.02, 0.19]	<b>-0.15 [-0.26, -0.05]</b>

text:  $P(\beta < 0) = 0.95$ ) but are better at recognizing machine-generated media (audio:  $P(\beta > 0) = 1.00$ ; image:  $P(\beta > 0) = 0.99$ ; text:  $P(\beta > 0) = 0.99$ ). Finally, FAM helps participants recognize human-generated audio data while hindering their ability to detect image and text data (audio:  $P(\beta > 0) = 1.00$ ; image:  $P(\beta < 0) = 1.00$ ; text:  $P(\beta < 0) = 1.00$ ). The results are flipped for machine-generated data (audio:  $P(\beta < 0) = 1.00$ ; image:  $P(\beta > 0) = 1.00$ ; text:  $P(\beta > 0) = 0.99$ ).

Additionally, we found several predictors that showed up only for some media types: We found support that holistic thinking (AHS) helped participants discern human-generated text data (text:  $P(\beta > 0) = 0.96$ ). The trend is reversed and stronger when classifying machine-generated data (audio:  $P(\beta > 0) = 0.98$ ; text:  $P(\beta < 0) = 0.97$ ). We also detected a general trend that holistic thinking worsened the participants’ performance on image data (both human- and machine-generated) (human:  $P(\beta < 0) = 0.92$ ; machine:  $P(\beta < 0) = 0.85$ ), but both distributions overlap with zero. When participants leaned more towards materialism (i.e., more conservative values), it worsened their performance in detecting machine-generated audio ( $P(\beta < 0) = 0.95$ ) and text ( $P(\beta < 0) = 0.98$ ) condition, as well as correctly classifying human-generated images ( $P(\beta < 0) = 1.0$ ). It did help participants to correctly classify machine-generated images ( $P(\beta > 0) = 1.0$ ).

Finally, when examining media literacy, we use the four subscales proposed in the original publication [106]: First, functional consuming (NMLS FC), i.e., being – technically – able to consume media, had a strong negative effect on detecting human-generated audio samples ( $P(\beta < 0) = 0.99$ ) and machine-generated image samples ( $P(\beta < 0) = 1.00$ ). We also detect a strong positive effect on correctly classifying images made by humans ( $P(\beta > 0) = 1.00$ ). Second, critical consumption (NMLS CC), i.e., the ability to conceive media messages as subjective instead of neutral, exhibits a positive effect for human-generated text data ( $P(\beta > 0) = 1.00$ ) and negative effects for both human-generated image data ( $P(\beta < 0) = 1.00$ ) and machine-generated text data ( $P(\beta < 0) = 1.00$ ). Third, critical prosuming (NMLS CP), i.e., the ability to critically analyze self-created work, has a positive effect on human-generated audio and text data (audio:  $P(\beta > 0) = 0.99$ ; text:  $P(\beta > 0) = 1.00$ ) and the results are

flipped for machine-generated data (audio:  $P(\beta < 0) = 1.00$ ; text:  $P(\beta < 0) = 1.00$ ). Finally, functional prosuming, i.e., the ability to technically construct media, only showed a positive effect on detecting machine-generated images media ( $P(\beta > 0) = 0.95$ ).

**Finding 5.** We found that GTS, CRT, FAM, and NMLS CP consistently showed up across all media types and AHS, NMLS FC, and NMLS CC for one or more media types. A key observation is that while they often have strong influences, they counteract each other, e.g., a higher CRT score predicts better performance for recognizing machine-generated media but at the same time hinders the ability to detect human-generated samples. We hypothesize that these observations match those in Section 4.3.3, i.e., that AI-generated media is already at a level where people are mainly guessing, thus the influence of these variables collapses. This highlights the need for more research into automatically detecting machine-generated media.

## 5. Discussion

With this work, we aim to gain a better understanding of the current state of people’s ability to detect AI-generated media in different countries. Overall, the results of our study indicate that machine-generated media has become almost indistinguishable from “real” media. This is especially true in the image domain, where participants from all countries performed even worse than random guessing.

This trend is not fully universal. We found indications that German audio data still lacks behind. Intuitively, one might suspect that Western media would be of higher quality, since it is more often used for research purposes. For example, the original training set for GPT-3 [53] (our text model) does not include much Chinese text (0.16%) when compared to English (93.69%) (number of documents).

While there was a statistical difference between how Chinese and US participants rated machine-generated media, we could not detect the same difference between accuracy ratings (cf. Section 4.1). When we analyzed this phenomenon

in more detail, we found a consistent pattern that Chinese and US participants seem to resort to random guessing (cf. Section 4.1 and Section 4.3). However, they did perform differently when looking at human- and machine-generated media separately, which explains the difference in ratings. This suggests that machine-generated English media is not that far ahead as typically believed in the machine learning community [113].

We can further connect our results to prior research by studying key predictors based on our literature review. The most prominent are CRT, FAM, and GTS, which all have been found to be significant across all media types: These findings extend previous works, which have found GTS and CRT to be significant factors when predicting the likelihood of accepting friend requests from machine-generated LinkedIn profiles and the sharing of political deepfakes, respectively. Our results indicate that these singular findings might generalize to multiple media types and should be closely studied in the future. Partly in line with related work, we found conservatism (in our case materialistic values) to negatively influence the detection of artificially generated text and audio files [94]. Participants holding more materialistic values were worse at identifying forged images.

Contrary to our preregistration, these findings did not show up along the lines of Western (traditionally associated with more analytic thinking) and Asian (more holistic thinking) cultures. Additionally, while AHS showed up as a significant predictor in our regression analysis, the results were somewhat inconclusive. This highlights the importance of more cross-cultural research in this area.

## 6. Limitations

In this study, we investigated the ability of humans to identify human- and machine-generated media types detached from a particular context. We believe that our results serve as a basis for people's general abilities to detect fake content across different media types and countries. However, context might help or hinder people in their detection abilities. We can only speculate that participants might perform worse at detecting AI-generated media in a real-world scenario, as they were told in this study that some media were fake and were instructed to identify human- and machine-generated media. Therefore, future work should take context into consideration and/or study detection in real-world scenarios. It has to be noted that the conclusions drawn from our results must take the lack of context into account.

Generally speaking, it is difficult for online panels to reach people with low education and older people (see [123], [124]). Thus, we were not able to fully meet all representative quotas, especially for China. However, as generated media are more likely to be displayed on the Internet, we believe our sample still meets the target group, and we considered representative shortcomings in our analysis (see Section 4.3). Nevertheless, we believe it would be worthwhile to sample especially the aforementioned underrepresented groups in future work.

Furthermore, we could not accurately account for control variables, as our panel provider could not provide us with accurate information on which devices our participants used. Additionally, the majority of participants in the audio group did not indicate which device they used for listening to the audio files. Thus, we abstain from entering these variables into the regression.

Finally, there is an expiration date to our conclusions. The field of generative modelling is one of the most active in the machine learning community. For example, during our study, diffusion models have almost completely replaced GAN-based methods. These models are often paired with large-language models, generating impressive art pieces from simple text prompts. While their use case is confined to generative artworks at the moment, one can suspect they will be used for malicious use cases in the future. In the same vein, there have been several projects in recent years that address the problem of underrepresented languages or speech corpora [125]–[127].

## 7. Conclusion

In this paper we present the first cross-country, cross-media study on people's ability to detect AI-generated media. We found that media generated with current state-of-the-art methods has become virtually indistinguishable from "real" media. Across all countries and all media types, people rated AI-generated samples as more likely to be produced by a human than a machine. Additionally, our statistical analysis showed that participants mostly guessed randomly, and it was very challenging for them to decide which stimuli were machine-generated and what clues to look for. Further, we found that generalized trust, cognitive reflection, familiarity, holistic thinking, political orientation, and some sub-scales of News Media Literacy had an effect on the ability to discern human- and machine-generated media.

Our results clearly show that machine-generative media are indistinguishable from real media. Since perfect technical detection seems unattainable, we argue that future research should not focus on how to avoid generative AI but rather, how to live with it. One relevant aspect is the inclusion of context: People might perceive machine-generated media differently when they see it in news articles, social media posts, or advertisements. In addition, research should explore which applications of AI are acceptable to humans and which are not. For example, are people okay with automatically generated summaries of product reviews? Generated news articles? Chatbots for therapy? We believe that it is the responsibility of policymakers to proactively regulate the use of generated media based on scientific evidence to take ethical and societal aspects into account. Only careful legislation that takes human perception and their values into account can mitigate the harmful effects of artificially generated media, without hindering its positive impacts. Whatever the future might hold, generative AI is here to stay, and we hope that this study can serve as a wake-up call for increased human-centered research into these methods.



## Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2092 CASA – 390781972 and the German Federal Ministry of Education and Research under the grants UbiTrans (16KIS1900) and AIGenCY (16KIS2012).

## References

- [1] Y. N. Harari, *Homo Deus: A brief History of Tomorrow*. Random House, 2016.
- [2] Z. Sharf, "'The Shining' Deepfake Goes Viral With Jim Carrey Starring as Jack Torrance," *IndieWire*, 2019.
- [3] C. Stupp, "Fraudsters Used AI to Mimic CEO's Voice in Unusual Cybercrime Case," *The Wall Street Journal*, 2019.
- [4] N. Thompson and I. Lapowsky, "How Russian Trolls Used Meme Warfare to Divide America," *Wired*, 2017.
- [5] K. Hao, "The Biggest Threat of Deepfakes isn't the Deepfakes Themselves," *MIT Technology Review*, 2019.
- [6] P. Mwai, "Tigray conflict: The fake UN diplomat and other misleading stories," *BBC Reality Check*, 2021.
- [7] T. A. C. D. F. R. Lab, "Inauthentic Instagram accounts with synthetic faces target Navalny protests," *Medium*, 2021.
- [8] J. Wakefield, "Deepfake presidents used in Russia-Ukraine war," *BBC Technology*, 2022.
- [9] X. Zhang, S. Karaman, and S.-F. Chang, "Detecting and Simulating Artifacts in GAN Fake Images," in *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2019.
- [10] Y. Qian, G. Yin, L. Sheng, Z. Chen, and J. Shao, "Thinking in Frequency: Face Forgery Detection by Mining Frequency-Aware Clues," in *European Conference on Computer Vision (ECCV)*, 2020.
- [11] N. Yu, L. S. Davis, and M. Fritz, "Attributing Fake Images to GANs: Learning and Analyzing GAN Fingerprints," in *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [12] J. Frank, T. Eisenhofer, L. Schönherr, A. Fischer, D. Kolossa, and T. Holz, "Leveraging Frequency Analysis for Deep Fake Image Recognition," in *International Conference on Machine Learning (ICML)*, 2020.
- [13] F. Marra, D. Gragnaniello, L. Verdoliva, and G. Poggi, "Do GANs Leave Artificial Fingerprints?" In *IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2019.
- [14] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, and A. A. Efros, "CNN-generated images are surprisingly easy to spot... for now," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [15] S. Tariq, S. Lee, H. Kim, Y. Shin, and S. S. Woo, "GAN is a Friend or Foe? A Framework to Detect Various Fake Face Images," in *ACM/SIGAPP Symposium on Applied Computing*, 2019.
- [16] S. McCloskey and M. Albright, "Detecting GAN-Generated Imagery Using Color Cues," *arXiv preprint arXiv:1812.08247*, 2018.
- [17] L. Nataraj, T. M. Mohammed, B. Manjunath, *et al.*, "Detecting GAN Generated Fake Images Using Co-Occurrence Matrices," *Electronic Imaging*, 2019.
- [18] H. Mo, B. Chen, and W. Luo, "Fake Faces Identification via Convolutional Neural Network," in *ACM Workshop on Information Hiding and Multimedia Security*, 2018.
- [19] F. Marra, D. Gragnaniello, D. Cozzolino, and L. Verdoliva, "Detection of GAN-Generated Fake Images over Social Networks," in *IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2018.
- [20] J. Frank and L. Schönherr, "WaveFake: A Data Set to Facilitate Audio Deepfake Detection," in *Conference on Neural Information Processing Systems - Datasets and Benchmarks Track*, 2021.
- [21] J. Ricker, S. Damm, T. Holz, and A. Fischer, "Towards the Detection of Diffusion Model Deepfakes," *arXiv preprint arXiv:2210.14571*, 2022.
- [22] J. Pu, Z. Sarwar, S. M. Abdullah, *et al.*, "Deepfake Text Detection: Limitations and Opportunities," in *IEEE Symposium on Security and Privacy (S&P)*, 2022.
- [23] N. Hulzebosch, S. Ibrahim, and M. Worring, "Detecting CNN-Generated Facial Images in Real-World Scenarios," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [24] Y. Hwang, J. Y. Ryu, and S.-H. Jeong, "Effects of Disinformation Using Deepfake: The Protective Effect of Media Literacy Education," *Cyberpsychology, Behavior, and Social Networking*, 2021.
- [25] J. Mink, L. Luo, N. M. Barbosa, O. Figueira, Y. Wang, and G. Wang, "DeepPhish: Understanding User Trust Towards Artificially Generated Profiles in Online Social Networks," in *USENIX Security Symposium*, 2022.
- [26] S. J. Nightingale and H. Farid, "AI-synthesized Faces Are Indistinguishable from Real Faces and More Trustworthy," *Proceedings of the National Academy of Sciences*, 2022.
- [27] F. Lago, C. Pasquini, R. Böhme, H. Dumont, V. Goffaux, and G. Boato, "More Real Than Real: A Study on Human Visual Perception of Synthetic Faces [Applications Corner]," *IEEE Signal Processing Magazine*, 2022.
- [28] N. M. Müller, K. Pizzi, and J. Williams, "Human Perception of Audio Deepfakes," in *Proceedings of the 1st International Workshop on Deepfake Detection for Audio Multimedia*, 2022.
- [29] M. Appel and F. Prielzel, "The detection of political deepfakes," *Journal of Computer-Mediated Communication*, 2022.
- [30] "Who inadvertently shares deepfakes? Analyzing the role of political interest, cognitive ability, and social network size," *Telematics and Informatics*, 2021.
- [31] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, *et al.*, "Generative Adversarial Networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2014.
- [32] A. Radford, L. Metz, and S. Chintala, "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks," in *International Conference on Learning Representations (ICLR)*, 2016.
- [33] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [34] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved Techniques for Training GANs," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [35] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein GAN," in *International Conference on Machine Learning (ICML)*, 2017.
- [36] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved Training of Wasserstein GANs," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [37] H. Petzka, A. Fischer, and D. Lukovnikov, "On the Regularization of Wasserstein GANs," in *International Conference on Learning Representations (ICLR)*, 2018.
- [38] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral Normalization for Generative Adversarial Networks," *International Conference on Learning Representations (ICLR)*, 2018.
- [39] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive Growing of GANs for Improved Quality, Stability, and Variation," in *International Conference on Learning Representations (ICLR)*, 2018.
- [40] T. Karras, S. Laine, and T. Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [41] A. Brock, J. Donahue, and K. Simonyan, "Large Scale GAN Training for High Fidelity Natural Image Synthesis," in *International Conference on Learning Representations (ICLR)*, 2019.
- [42] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and Improving the Image Quality of StyleGAN," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [43] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.



- [44] A. Q. Nichol and P. Dhariwal, "Improved Denoising Diffusion Probabilistic Models," in *International Conference on Machine Learning (ICML)*, 2021.
- [45] P. Dhariwal and A. Nichol, "Diffusion Models Beat GANs on Image Synthesis," 2021.
- [46] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," 2021.
- [47] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical Text-Conditional Image Generation with CLIP Latents," *arXiv preprint arXiv:2204.06125*, 2022.
- [48] D. Bahdanau, K. Cho, and Y. Bengio, "Neural Machine Translation by Jointly Learning to Align and Translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [49] K. Cho, B. Van Merriënboer, C. Gulcehre, et al., "Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [50] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to Sequence Learning with Neural Networks," 2014.
- [51] Y. Kim, C. Denton, L. Hoang, and A. M. Rush, "Structured Attention Networks," 2017.
- [52] A. Vaswani, N. Shazeer, N. Parmar, et al., "Attention is all you need," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [53] T. B. Brown, B. Mann, N. Ryder, et al., "Language Models are Few-Shot Learners," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [54] OpenAI, *GPT-4 Technical Report*, 2023. arXiv: 2303.08774 [cs.CL].
- [55] A. Chowdhery, S. Narang, J. Devlin, et al., "PaLM: Scaling Language Modeling with Pathways," *arXiv preprint arXiv:2204.02311*, 2022.
- [56] H. Zen, K. Tokuda, and A. W. Black, "Statistical Parametric Speech Synthesis," *Speech Communication*, 2009.
- [57] K. Tokuda, T. Yoshimura, T. Masuko, T. Kobayashi, and T. Kitamura, "Speech Parameter Generation Algorithms for HMM-Based Speech Synthesis," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2000.
- [58] T. Yoshimura, K. Tokuda, T. Masuko, T. Kobayashi, and T. Kitamura, "Simultaneous Modeling of Spectrum, Pitch and Duration in HMM-Based Speech Synthesis," in *Sixth European Conference on Speech Communication and Technology*, 1999.
- [59] K. Kumar, R. Kumar, T. de Boissiere, et al., "MelGAN: Generative Adversarial Networks for Conditional Waveform Synthesis," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [60] M. Bińkowski, J. Donahue, S. Dieleman, et al., "High Fidelity Speech Synthesis with Adversarial Networks," in *International Conference on Learning Representations (ICLR)*, 2020.
- [61] C. Donahue, J. McAuley, and M. Puckette, "Adversarial Audio Synthesis," in *International Conference on Learning Representations (ICLR)*, 2019.
- [62] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative Adversarial Networks for Efficient and High Fidelity Speech Synthesis," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [63] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel WaveGAN: A Fast Waveform Generation Model Based on Generative Adversarial Networks with Multi-Resolution Spectrogram," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- [64] V. L. Rubin, "Disinformation and Misinformation Triangle: A Conceptual Model for "Fake News" Epidemic, Causal Factors and Interventions," *Journal of Documentation*, 2019.
- [65] S. M. Jang and J. K. Kim, "Third Person Effects of Fake News: Fake News Regulation and Media Literacy Interventions," *Computers in Human Behavior*, 2018.
- [66] S. M. Jones-Jang, T. Mortensen, and J. Liu, "Does Media Literacy Help Identification of Fake News? Information Literacy Helps, but Other Literacies Don't," *American Behavioral Scientist*, vol. 65, 2021.
- [67] S.-H. Jeong, H. Cho, and Y. Hwang, "Media Literacy Interventions: A Meta-Analytic Review," *Journal of Communication*, 2012.
- [68] D. J. Munro, "Individualism and Holism: Studies in Confucian and Taoist Values," *Michigan Monographs in Chinese Studies*, 1985.
- [69] H. Nakamura, *Ways of Thinking of Eastern Peoples: India, China, Tibet, Japan*. Motilal Banarsidass Publishes, 1991.
- [70] T. Masuda and R. E. Nisbett, "Attending holistically versus analytically: comparing the context sensitivity of Japanese and Americans," *Journal of Personality and Social Psychology*, 2001.
- [71] R. E. Nisbett, K. Peng, I. Choi, and A. Norenzayan, "Culture and Systems of Thought: Holistic Versus Analytic Cognition," *Psychological Review*, 2001.
- [72] G. Pennycook and D. G. Rand, "Who Falls for Fake News? The Roles of Bullshit Receptivity, Overclaiming, Familiarity, and Analytic Thinking," *Journal of Personality*, 2020.
- [73] M. V. Bronstein, G. Pennycook, A. Bear, D. G. Rand, and T. D. Cannon, "Belief in Fake News is Associated with Delusionality, Dogmatism, Religious Fundamentalism, and Reduced Analytic Thinking," *Journal of applied research in memory and cognition*, 2019.
- [74] J. F. Helliwell, "Well-Being, Social Capital and Public Policy: What's New?" *The Economic Journal*, 2006.
- [75] R. D. Putnam, "Bowling Alone: America's Declining Social Capital," in *Culture and Politics*, 2000.
- [76] R. C. Mayer, J. H. Davis, and F. D. Schoorman, "An integrative Model of Organizational Trust," *Academy of Management Review*, 1995.
- [77] W. W. Maddux and M. B. Brewer, "Gender Differences in the Relational and Collective Bases for Trust," *Group Processes & Intergroup Relations*, 2005.
- [78] N. Ashraf, I. Bohnet, and N. Piankov, "Decomposing Trust and Trustworthiness," *Experimental Economics*, 2006.
- [79] B. Rothstein and E. M. Uslaner, "All for all: Equality, corruption, and social trust," *World Politics*, 2005.
- [80] W. Tov and E. Diener, "The Well-Being of Nations: Linking Together Trust, Cooperation, and Democracy," in *The Science of Well-Being*, 2009.
- [81] J. L. Gibbs, N. B. Ellison, and C.-H. Lai, "First Comes Love, Then Comes Google: An Investigation of Uncertainty Reduction Strategies and Self-Disclosure in Online Dating," *Communication Research*, 2011.
- [82] X. Ma, E. Sun, and M. Naaman, "What Happens in Happn: The Warranting Powers of Location History in Online Dating," in *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, 2017.
- [83] E. Ert, A. Fleischer, and N. Magen, "Trust and Reputation in the Sharing Economy: The Role of Personal Photos in Airbnb," *Tourism Management*, 2016.
- [84] A. Lampinen and C. Cheshire, "Hosting via Airbnb: Motivations and Financial Assurances in Monetized Network Hospitality," in *CHI Conference on Human Factors in Computing Systems*, 2016.
- [85] K. E. Stanovich and R. F. West, "Individual differences in reasoning: Implications for the rationality debate?" *Behavioral and Brain Sciences*, 2000.
- [86] S. Frederick, "Cognitive Reflection and Decision Making," *Journal of Economic Perspectives*, 2005.
- [87] G. Pennycook and D. G. Rand, "Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning," *Cognition*, 2019.
- [88] B. Bago, D. Rand, and G. Pennycook, "Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines," *Journal of Experimental Psychology: General*, 2020.
- [89] G. Pennycook and D. G. Rand, "The Psychology of Fake News," *Trends in Cognitive Sciences*, 2021.
- [90] C. Batailler, S. M. Brannon, P. E. Teas, and B. Gawronski, "A Signal Detection Approach to Understanding the Identification of Fake News," *Perspectives on Psychological Science*, 2022.
- [91] M. Mosleh, G. Pennycook, A. A. Arechar, and D. G. Rand, "Cognitive reflection correlates with behavior on Twitter," *Nature Communications*, 2021.

- [92] M. Groh, A. Sankaranarayanan, A. Lippman, and R. Picard, *Human Detection of Political Deepfakes across Transcripts, Audio, and Video*, 2022.
- [93] S. Ahmed, "Fooled by the fakes: Cognitive differences in perceived claim accuracy and sharing intention of non-political deepfakes," *Personality and Individual Differences*, 2021.
- [94] D. P. Calvillo, B. J. Ross, R. J. B. Garcia, T. J. Smelter, and A. M. Rutchick, "Political Ideology Predicts Perceptions of the Threat of COVID-19 (and Susceptibility to Fake News About It)," *Social Psychological and Personality Science*, 2020.
- [95] A. Chadwick and C. Vaccari, *News Sharing on UK Social Media Misinformation, Disinformation, And Correction*. Loughborough University, 2019.
- [96] J. Shen, R. Pang, R. J. Weiss, et al., "Natural TTS Synthesis by Conditioning WaveNet on Mel Spectrogram Predictions," in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [97] K. Ito and L. Johnson, *The LJ Speech Dataset*, <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [98] L. Databaker (Beijing) Technology Co., *Chinese Standard Mandarin Speech Copus*, 2017.
- [99] P. Puchler, J. Wirth, and R. Peinl, "HUI-Audio-Corpus-German: A high quality TTS dataset," in *German Conference on Artificial Intelligence (Künstliche Intelligenz)*, 2021.
- [100] O. Parkhi, A. Vedaldi, and A. Zisserman, "Deep Face Recognition," 2015.
- [101] *National Public Radio*. Accessed: December 13, 2023.
- [102] *Tagesschau*. Accessed: December 13, 2023.
- [103] *China Central Television*. Accessed: December 13, 2023.
- [104] *AllSlides*. Accessed: December 13, 2023.
- [105] *OpenAI*. Accessed: December 13, 2023.
- [106] M. Koc and E. Barut, "Development and Validation of New Media Literacy Scale (NMLS) for University Students," *Computers in Human Behavior*, 2016.
- [107] M. Martín-Fernández, B. Requero, X. Zhou, D. Gonçalves, and D. Santos, "Refinement of the Analysis-Holism Scale: A Cross-Cultural Adaptation and Validation of Two Shortened Measures of Analytic Versus Holistic Thinking in Spain and the United States," *Personality and Individual Differences*, 2022.
- [108] T. Yamagishi and M. Yamagishi, "Trust and Commitment in the United States and Japan," *Motivation and emotion*, 1994.
- [109] R. Inglehart and P. R. Abramson, "Measuring Postmaterialism," *American Political Science Review*, 1999.
- [110] International Chamber of Commerce and European Society for Opinion and Marketing Research, *International Code on Market and Social Research*, 2007.
- [111] J. Pater, A. Coupe, R. Pfafman, C. Phelan, T. Toscos, and M. L. Jacobs, "Standardizing reporting of participant compensation in HCI: A systematic literature review and recommendations for the field," in *ACM Conference on Human Factors in Computing Systems*, ser. CHI '21, Yokohama, Japan: Association for Computing Machinery, 2021.
- [112] A. W. Meade and S. B. Craig, "Identifying Careless Responses in Survey Data.," *Psychological Methods*, 2012.
- [113] J. W. Rae, S. Borgeaud, T. Cai, et al., "Scaling Language Models: Methods, Analysis & Insights from Training Gopher," *arXiv preprint arXiv:2112.11446*, 2021.
- [114] D. Deffner, J. M. Rohrer, and R. McElreath, "A Causal Framework for Cross-Cultural Generalizability," *Advances in Methods and Practices in Psychological Science*, 2022.
- [115] H. Goldstein, *Multilevel Statistical Models*. John Wiley & Sons, 2011.
- [116] A. Gelman and J. Hill, *Data Analysis using Regression and Multilevel/Hierarchical Models*. Cambridge University Press, 2006.
- [117] S. Rabe-Hesketh and A. Skrondal, "Multilevel Modelling of Complex Survey Data," in *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2006.
- [118] D. Stegmüller, "How many countries for multilevel modeling? a comparison of frequentist and bayesian approaches," in *American Journal of Political Science*, 2013.
- [119] A. Vehtari, A. Gelman, and J. Gabry, "Practical Bayesian Model Evaluation using Leave-One-Out Cross-Validation and WAIC," *Statistics and Computing*, 2017.
- [120] R. McElreath, *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. Chapman and Hall/CRC, 2020.
- [121] A. Gelman, J. Hill, and M. Yajima, "Why we (usually) don't have to worry about multiple comparisons," *Journal of Research on Educational Effectiveness*, vol. 5, no. 2, pp. 189–211, 2012.
- [122] J. Pearl, "Comment: Understanding Simpson's Paradox," in *Probabilistic and Causal Inference: The works of Judea Pearl*, 2022, pp. 399–412.
- [123] E. M. Redmiles, S. Kross, and M. L. Mazurek, "How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples," in *IEEE Symposium on Security and Privacy (S&P)*, 2019.
- [124] J. Tang, E. Birrell, and A. Lerner, "Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys," in *Symposium on Usable Privacy and Security*.
- [125] P. Rust, J. Pfeiffer, I. Vulic, S. Ruder, and I. Gurevych, "How Good is Your Tokenizer? On the Monolingual Performance of Multilingual Language Models," in *Annual Meeting of the Association for Computational Linguistics (ACL/IJCNLP)*, C. Zong, F. Xia, W. Li, and R. Navigli, Eds., 2021.
- [126] T. Limisiewicz, D. Malkin, and G. Stanovsky, "You Can Have Your Data and Balance It Too: Towards Balanced and Efficient Multilingual Models," *CoRR*, 2022.
- [127] Y. Zhang, R. J. Weiss, H. Zen, et al., "Learning to Speak Fluently in a Foreign Language: Multilingual Speech Synthesis and Cross-Language Voice Cloning," in *Proceedings of Interspeech (INTERSPEECH)*, 2019.
- [128] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, "Probabilistic programming in python using PyMC3," *PeerJ Computer Science*, 2016.
- [129] *Aesara*. [Online]. Available: <https://github.com/aesara-devs/aesara>.
- [130] J. Bradbury, R. Frostig, P. Hawkins, et al., *JAX: Composable transformations of Python+NumPy programs*, 2018. [Online]. Available: <http://github.com/google/jax>.

## Appendix A. Statistical Analysis

We model the number of times  $Y_i$  a participant  $i$  correctly judged a media out of all his trials  $N_i$ , as a binomial distribution:

$$Y_i \sim \text{Binomial}(N_i, p_i) ,$$

where we parameterize the probability  $p_i$  by several different models. In our notion we use superscripts as names and subscripts as indexes; for example, in the expression  $\delta_{c,a}^{age}$  the parameter is named  $\delta^{age}$  and indexed by country  $c$  and age  $a$ . Our full-model (Age/Education/Time by Country) models the probability  $p_i$  as:

$$\text{logit}(p_i) = \alpha_c^{country} + \delta_{c,e}^{edu} + \delta_{c,a}^{age} + \delta_c^{time} X_i^{time}$$

$$\alpha_c^{country} \sim N(\mu^{country}, \sigma^{country})$$

$$\delta_{c,e}^{edu} \sim N(\mu^{edu}, \sigma^{edu})$$

$$\delta_c^{time} \sim N(\mu^{time}, \sigma^{time})$$

$$\mu^{country}, \mu^{edu}, \mu^{time} \sim N(0, 1.5)$$

$$\sigma^{country}, \sigma^{edu}, \sigma^{time} \sim \lambda(1)$$

$$\begin{pmatrix} \delta_{c,1}^{age} \\ \delta_{c,2}^{age} \\ \dots \\ \delta_{c,A}^{age} \end{pmatrix} \sim N \left[ \begin{pmatrix} 0 \\ 0 \\ \dots \\ 0 \end{pmatrix}, K \right]$$

$$K_{x,y} = \eta_c^2 \exp(-\rho_c^2 D_{x,y}^2) + \pi_{x,y} \sigma_c^2$$

$$\eta_c^2 \sim \lambda(\eta^{age})$$

$$\sigma_c^2 \sim \sigma^{age}$$

$$\rho_c^2 = \frac{1}{2} l_c^2$$

$$l_c \sim \text{HalfNormal}(l^{age})$$

$$\eta^{age} \sim \lambda(2)$$

$$l^{age} \sim \text{HalfNormal}(5)$$

$$\sigma^{age} \sim \lambda(1)$$


Intuitively, the model assumes that while each country has some freedom in the effect of the education/age levels, they should correlate across countries. Thus, our model uses partial-pooling to model these correlations. More specifically, the varying intercept  $\alpha_c^{country}$  per country is used to model general differences between the countries (e.g., quality of different generative models). We include two correction terms: the term  $\delta_{c,e}^{edu}$  for each education level per country and the term  $\delta_c^{time} X_i^{time}$  for the median time taken in the survey. When modelling the age of our participants, we would expect that similar age brackets (e.g., people around the ages 25-30 and 30-35) are characterized by more similar behavior than

age brackets that are far apart (e.g., 20-25 and 60-65). We divide the age of our participants into 17 categories, spanning 5 years each, and used a Gaussian process to model the interactions between the different levels with an age-specific offset  $\delta_{c,a}^{age}$ . The offset  $\delta_{c,a}^{age}$  is drawn from a multivariate Normal distribution with mean zero and a covariance matrix defined by the kernel function  $K$ . The covariance between any age pair  $x, y$  equals the maximum covariance  $\eta^2$ , which is reduced at the rate  $\rho^2$  by the squared distance between the two,  $D_{x,y}^2$ . There is an additional covariance parameter  $\sigma^2$  with a dummy variable  $\pi_{x,y}$  which indicates that  $x = y$ . Thus, it expresses the additional covariance within each age group. We include our predictor variables as fixed terms and model them either as continuous or ordinal variables [120]. During the analysis we standardize continuous scores and encode the Inglehart index, the CRT and knowledge score as ordinal variables.

All models were developed using pymc 4.3.0 [128] with aesara 2.8.7 [129] and jax 0.3.24 [130] as the backend. For each model, we started development on synthetic data, using prior predictive checks. When using real data, we fit the parameters of all models with MCMC using the NUTS sampler for 4,000 iterations (2,000 warm-up). Model fit was confirmed by inspecting  $\hat{r}$ -values. All analyses were run on a server running Ubuntu 18.04.6 with an Intel Xeon Gold 6130 CPU and 128GB RAM.

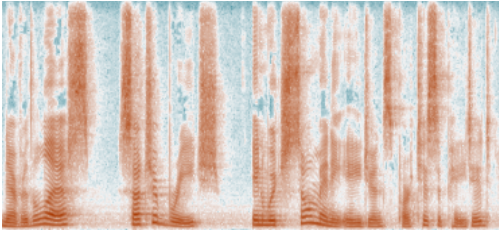
## Appendix B. Survey Design

Below we provide a screenshot illustrating our survey design (exemplary for the image condition).



## Appendix C. Stimuli

Below we report exemplary stimuli for the three types of media used in the study. For speech, we show the audio files in the frequency domain. The full set of all stimuli is shared together with the preregistration ([osf.io/xy6v5](https://osf.io/xy6v5)).

	Human-Generated	Machine-Generated
Image		
Audio		
Text	<p>Russia escalated its bombardment of the Ukrainian capital and launched new assaults on the port city of Mariupol, making bloody advances on the ground as Ukraine's president prepared Wednesday to make a direct appeal for more help in a rare speech by a foreign leader to the U.S. Congress. As the invasion entered its third week, Ukrainian President Volodymyr Zelenskyy suggested there was still some reason to be optimistic negotiations might yet yield an agreement with the Russian government.</p>	<p>A magnitude 7.3 earthquake struck off the coast of Fukushima on Wednesday evening, triggering a tsunami advisory and cutting power to more than 2 million Tokyo homes, authorities said. The quake knocked out power to about 1.9 million households in the capital region shortly after 6 p.m., NHK national television reported. That figure was later revised to just over 2 million households as crews repaired damage caused by an earlier blackout, it said.</p>