

# Casual Conversations v2: Designing a large consent-driven dataset to measure algorithmic bias and robustness

CANER HAZIRBAS, Meta AI, USA

YEJIN BANG\* and TIEZHENG YU\*, Hong Kong University of Science and Technology, Hong Kong

PARISA ASSAR, Meta AI, USA

BILAL PORGALI, Meta AI, United Kingdom

VÍTOR ALBIERO, Meta AI, USA

STEFAN HERMANEK, Meta AI, USA

JACQUELINE PAN, Meta AI, USA

EMILY MCREYNOLDS, Meta AI, USA

MIRANDA BOGEN, Meta AI, USA

PASCALE FUNG, Hong Kong University of Science and Technology, Hong Kong

CRISTIAN CANTON FERRER, Meta AI, USA

Developing robust and fair AI systems requires datasets with comprehensive set of labels that can help ensure the validity and legitimacy of relevant measurements. Recent efforts, therefore, focus on collecting person-related datasets that have carefully selected labels, including sensitive characteristics, and consent forms in place to use those attributes for model testing and development. Responsible data collection involves several stages, including but not limited to determining use-case scenarios, selecting categories (annotations) such that the data are fit for the purpose of measuring algorithmic bias for subgroups and most importantly ensure that the selected categories/subcategories are robust to regional diversities and inclusive of as many subgroups as possible.

Meta, in a continuation of our efforts to measure AI algorithmic bias and robustness<sup>1</sup>, is working on collecting a large consent-driven dataset with a comprehensive list of categories. This paper describes our proposed design of such categories and subcategories for *Casual Conversations v2*.

CCS Concepts: • **Computing methodologies** → **Computer vision; Natural language processing.**

Additional Key Words and Phrases: datasets, robustness, fairness, algorithmic bias

## ACM Reference Format:

Caner Hazirbas, Yejin Bang, Tiezheng Yu, Parisa Assar, Bilal Porgali, Vítor Albiero, Stefan Hermanek, Jacqueline Pan, Emily McReynolds, Miranda Bogen, Pascale Fung, and Cristian Canton Ferrer. 2022. Casual Conversations v2: Designing a large consent-driven dataset to measure algorithmic bias and robustness. In . ACM, New York, NY, USA, 18 pages.

---

\*Both authors contributed equally to this research.

<sup>1</sup><https://ai.facebook.com/blog/shedding-light-on-fairness-in-ai-with-a-new-data-set>

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

Manuscript submitted to ACM

## 1 INTRODUCTION

There is no shortage of literature [61, 92, 130] on AI systems performing inconsistently across different sub-groups of society, and scholars have continued to advocate for further diversity of such datasets. Several recent studies [8, 41, 55, 67, 75] propose various learning strategies for AI models to be well-calibrated across all protected subgroups, while others focus on collecting responsible datasets [57, 82, 124] to make sure evaluations of AI models are accurate and algorithmic bias can be measured while promoting data privacy.

There has been much criticism regarding the design choice of the publicly used datasets, such as for ImageNet [36, 38, 56, 70]. Discussions are mostly focused on concerns around collecting sensitive data about people without their consent. Casual Conversations v1 [57] was one of the first benchmarks that was designed with permission from participants. However, that dataset has several limitations: samples were collected only in the US, the gender label is limited to three options, and only age and gender labels are self-provided with the permission of the participants. While it has enabled many researchers [40, 51] to evaluate their models across these and other dimensions, it does not cover many categories that may be needed to more fully evaluate models such as physical attributes and more granular subcategories.

In order to understand the landscape of a responsible dataset for model evaluation in audio/vision/speech models, we conducted on a literature survey that helped us understand how we should identify new self-provided label options and how those labels should be sub-categorized in order to be more inclusive. The objective of the design is to provide a list of categories to measure potential algorithmic failure on as many use-cases as possible, *e.g.*, computer vision, audio, speech and language models. Therefore to facilitate this, we decided to collect ten categories. Many of the recent works [5–7, 32, 35, 51, 53, 57, 62, 72, 82, 92, 93, 115] break down their analysis mostly based on age, gender, race, language, geo-location, apparent skin-tone, and voice timbre. However, there are several other categories that have not been covered in most of the datasets but will improve the inclusion of less represented sub-groups of people and circumstances such as physical attributes (body size, body/face markings, tattoo, *etc.*), and activity (standing, walking, *etc.*).

While we do believe demography-related labels are highly salient to measuring bias, we also add a set of categories that are relevant to model performance for the relevant use-cases and thus can help measure robustness of the models that may directly impact or be experienced as model fairness gaps, such as different recording setups (lighting, camera quality, environment, *etc.*).

In Section 2 we provide our literature review for each category and present proposed subcategories of each that will be included in our dataset. Further, in Section 3 we discuss other categories that are commonly used for measuring algorithmic bias, *e.g.*, race, and elaborate the reasons why we opted not to solicit them in our data collection.

## 2 PROPOSED CATEGORIES IN THE CASUAL CONVERSATIONS V2 DATASET

In the following section, we discuss each of these categories, provide a list of proposed subcategories and further evaluate their utility for algorithmic bias measurement. In Table 1, we list all categories and indicate which ones will be self-provided and which ones will be annotated.

### 2.1 Age

Investigating the robustness of AI models for age groups is important in building fair AI algorithms. Notably, many current public datasets [57, 74] perform analysis on the various age groups. However the definition of these age groups vary from dataset to dataset and we would like to propose the categorization shown in Table 2. This proposal summarizes

Section	Category	Self-provided	Annotated
Section 2.1	Age	✓	
Section 2.2	Gender	✓	
Section 2.3	Language/Dialect	✓	
Section 2.4	Geo-Location (Country/State/City)	✓	
Section 2.5	Disability	✓	
Section 2.6	Physical Adornments and Physical Attributes	✓	
Section 2.7	Voice Timbre		✓
Section 2.8	Apparent Skin Tone		✓
Section 2.9	Different Recording Setup		✓
Section 2.10	Activity		✓
Section	Excluded Categories		
Section 3	Race & Ethnicity	✗	✗
Section 3	Facial Expressions	✗	✗

Table 1. List of proposed self-provided and annotated categories that will present in the CCv2 dataset. We exclude several categories in this dataset and provide reasons in Section 3.

how we plan to categorize age, which is synthesized from various governmental and academic resources. Though fairness for AI with regard to youth is a priority for policymakers, we propose to collect data from people who are older than the age of maturity, which varies by country between (18 – 22) given equally salient privacy concerns when it comes to younger populations [49].

**Governmental resources.** Many government population statistics include age distribution. Here we consult the statistical standards of age groups from multiple governmental resources. Both Singapore Department of Statistics [106] and World Bank [16] divide the age groups into 0 – 14, 15 – 64, and 65+. Statistics Canada [29] chooses to divide the age groups into Children (0 – 14), Youth (15 – 24), Adults (25 – 64) and Seniors (65+). In addition, United Nations [101] uses 0 – 14, 15 – 24, 25 – 44, 45 – 64 and 65+ age groups.

**Scientific resources.** We also reviewed the current public datasets. Many works have been performed on using face representation to predict age groups [1, 39, 42, 64]. For example, an age group classification dataset based on facial features [60] divides the groups into Baby (0 – 2), Young Adults (3 – 39), Middle-age Adults (40 – 59), and Old Adults (60+). An Automatic Speech Recognition dataset that aims to quantify age bias [48] introduces age groups includes children (7 – 11), Teenagers (12 – 16), Adults (17 – 64) and Older Adults (65+). Casual Conversations v1 dataset [57] groups the category for age into three buckets that are 18 – 30, 31 – 45 and 46 – 85. FairFace: a face attribute dataset [69] and Common Voice: a speech corpus [13] use age buckets mainly at ten-year intervals.

Age Groups	Age Intervals
Young Adults	X-24
Adults	25-40
Middle-age Adults	41-65
Old Adults	65+

Table 2. The proposed age intervals for age groups. X, the age of maturity varies by country between 18 – 22.

## 2.2 Gender

Gender is one of the most actively discussed protected attributes when it comes to algorithmic bias and robustness of AI systems [26, 129, 136, 137, 144]. Until recent years and even today in many cultures, gender identities have often been interchangeably confused with sex at birth (male and female). Thus, cumulative data or data annotation in much research considers gender as binary (man and woman). Most of the facial recognition research [35, 53, 62, 115] only considers binary genders in their analyses. However, gender is a social construction that changes over time and differs among various cultures [110]. Thus, choosing subcategories in gender requires a thorough assessment.

There is no fixed number of subgroups for gender identity; rather, it is more of a spectrum than a discrete matter [25]. There are more than *twenty five* gender subcategories according to the LGBTQIA Resource Center Glossary [30]. However, as stated in [66] (guideline for gender options in the surveys), an exhaustive list with every possible option may overwhelm and cause fatigue to the participants. In order to ensure inclusion of all possible gender identities, we recommend offering “*Not stated above, please specify*” option where participants can indicate their genders in a free-form text. The following provides the overview of gender categorization in perspectives from governmental, scientific and survey resources.

**Governmental resources.** Conventionally, many governments have adopted heteronormative binary gender, *i.e.*, “*Cis Man/Male (M)*” and “*Cis Woman/Female (F)*”, as a common choice for legal and institutional systems and official documents. However, with increased awareness of “*Non-binary*” genders in past few years, there has been legal recognition with an option of “*Non-binary*” [33, 104, 108]. Seventeen countries around the world legally recognize the *Non-binary* gender option on their passports (*i.e.*, gender-neutral passport), including Argentina, Austria, Australia, Canada, Colombia, Denmark, Germany, Iceland, India, Ireland, Malta, Nepal, the Netherlands, New Zealand, Pakistan and the United States<sup>2</sup> [44]. Besides, there is sometimes an option not to disclose their gender identity (*e.g.*, “*Preferred not to say*”) in some governmental survey or consensus, respecting one’s willingness for the disclosure. However, further specification about the *Non-binary* genders is still very limited.

In the United States, many official documents now require the option of “*Non-binary*” as an effort to promote greater equity and inclusion. This includes adding “*Non-binary*” as a gender marker printed on the U.S. passport (starting from 2023) [104] and in the Equal Employment Opportunity Commission (EEOC)’s charge intake process [33]. A survey of the LGBTQ+ population by the UK government [108] has given more options including “*Trans Man/Woman*”, “*I do not want to say*”, “*I do not know*” and “*Something else*”. The “*Something else*” option allows participants of the survey to indicate their gender by themselves. The Australian government releases a standard for sex, gender, variations of sex characteristics and sexual orientation variables to standardize the collection and dissemination of related data [105]. Table 3 shows the considered or required gender options for aforementioned consensus or surveys.

**Scientific resources.** Gender is one of the most important protected attributes discussed within the topics of AI fairness. Broadly speaking, there are two cases where gender categorization is required 1) perceived gender: labeling the data with different gender categories when the data already includes gendered information (*e.g.*, utilizing gender pronouns for constructing data for gender-bias in text or labelling human figures for gender in images) 2) gender identity of participants: collecting the personal information from data construction contributors (*e.g.*, gender of actors in video data, gender of data annotators, gender of the authors of certain online comments). The former are cases where gender information has not been provided directly thus a third party assumes the gender information by observing

<sup>2</sup>The data are as of September 2022.

Considered Genders	UK National LGBT Survey [108]	Gender Survey Suggestion [66]	U.S. Passport [104]	U.S. EEOC [33]	Australian Govt. [105]
Cis Woman	Woman/girl	Female	Female (F)	Female (F)	Female (Sex at birth)
Cis Man	Man/Boy	Male	Male (M)	Male (M)	Male (Sex at birth)
Trans Man	Trans Man	Transgender Male			
Trans Woman	Trans Woman	Transgender Female			
Non-binary	Non-binary/ Genderqueer/ Agender/ Gender-fluid		Another gender identity (X)	Non-binary (X)	Non-binary
Not stated above (free-form)	I do not know	Gender Variant / Not-Conforming			[I/they] use a different term (please specify)
Prefer Not to Say	Something else I do not want to say	Not Listed Prefer Not to Answer			Prefer not to answer

Table 3. Gender subcategories considered or/and required by different governmental resources and social perspectives. EEOC stands for the Equal Employment Opportunity Commission.

already existing data. The latter, in contrast, is where gender information can be gathered from the participants or the data builder. We looked into both approaches to gender categorizations in scientific resources.

Although there has been increased awareness regarding gender inclusiveness in the AI research community, binary gender is still a common categorization choice of gender attribute [48, 77, 90, 99, 112, 126]. For instance, there are only a few exceptions incorporating “*Non-binary*” genders despite of the extensive list of research on gender bias mitigation in NLP community in past few years [129]. One of the reasons for considering binary genders is the difficulty of data collection with “*Non-binary*” genders. For example, gendered pronouns are clearer for binary genders (e.g., she, he) when crawling data from web-data in the NLP research. As a result, there is scarcity of such data that includes “*Non-binary*” genders (nonbinary people may also be referred to with pronouns they do not prefer in written context, and a common pronoun, “they,” also more commonly refers to groups of people). As an effort for inclusion, [90, 136, 137, 144] showed understanding of “*Non-binary*” genders and gender fluidity, but mostly limited to the acknowledgement. Some scholars go further and provide rationale on why “*Non-binary*” genders are not included to compensate their limited inclusion of gender diversity.

“*Non-binary*” genders are included when the gender of the data building contributors can be explicitly asked. Maarten *et al.* [121] included “*Non-binary*” genders option for collecting the personal information of data annotators and commonly used speech data set Common Voice [13] extended the options to have “*Other*”.

**Research Guideline/Survey Design Guideline.** While the topic of inclusion and diversity regarding gender has recently received more attention in AI and machine learning research, other research fields have extensive explorations that can be brought into AI/ML research. There are inclusive language guidelines for asking gender when collecting data from an individual from various institutions or organizations [66]. In addition to the binary genders (Woman and Man), most commonly recommended options are “*Non-binary*” and “*Non-disclosure*” option (e.g., “*Prefer not to answer/say/respond/disclose*”). Many of the guidelines also recommend to include “*Transgender(s)*” as an option besides “*Non-binary*”, e.g., some include it as “*Transgender*” [91] while some further separate “*Trans Man*” and “*Trans Woman*” [107]. An option of additional response (self-indication) is also recommended by most of the guidelines. To provide the self-indication option, it is recommended to use a term of “*Not Listed/Referred*” to avoid any negative connotation of the term “*Other*” which is a common practice for indicating the not-listed options [66].

**Proposed final list of subcategories.** While we avoid exhaustion-of-choices and survey fatigue, we came to a final list of subcategories as follow: *Cis Man*, *Cis Woman*, *Transgender Man*, *Transgender Woman*, *Non-binary*, *Not stated above, please specify (free-form)*, *Preferred not to say*. We also provide definitions of these gender terms to the participants during data collection:

- *Cisgender*: an individual whose gender identity aligns with those typically associated with the biological sex assigned to them at birth.
- *Transgender* (also shortened to trans) refers to an individual whose gender identity and/or expression is different from cultural expectations based on the sex they were assigned at birth. Being transgender does not imply any sexual orientation.
  - *Transgender man*: an individual whose biological sex was assigned female at birth, but identifies as a man.
  - *Transgender woman*: an individual whose biological sex was assigned male at birth, but identifies as a woman.
- *Non-binary*: a term used to describe a person who does not identify exclusively as a man or a woman. Non-binary people can identify as being both a man and a woman, somewhere in between, or fall completely outside the gender binary. Non-binary can also be used as an umbrella term encompassing identities such as genderqueer or gender-fluid.

### 2.3 Language/Dialect

Linguists have long established categories and subcategories for languages and dialects. There is likely to be no need for AI researchers to re-define this categorization for languages. Rather, it is important to understand which languages to include when constructing a multi-lingual dataset. We conducted our literature review in two directions, *i.e.*, 1) most commonly used languages and 2) a more diverse set of languages that takes into account low resource languages [114].

When we enriched the selected languages/dialects to a more diverse set, we gathered more than a *hundred* subcategories. In order to find the right balance, we first listed the most commonly spoken/used *nineteen* languages and extended them to a more diverse set with several other languages. These *nineteen* languages include English, Spanish, Japanese, Mandarin/Cantonese/Wu (Shanghainese), Korean, Vietnamese, Hindi/Tamil/Telugu, Arabic, Farsi, French, German, Italian, Slavic Languages, Portuguese, Turkish, Swahili, Thai, Bahasa (Indonesian) and American Sign Language (ASL). The additional more commonly used languages are Dutch, Greek, Swedish, Finnish, Hungarian, Filipino, Hebrew, and Bengali.

Languages also have linguistic variations (*i.e.*, dialects) that differ in grammar, pronunciation, and/or vocabulary [132, 133]. In practice of collecting a dataset, dialects are often not clearly indicated, either because they are not considered or in some cases are included as optional. In the online crowd-sourced speech data Common Voice [13] dataset, volunteers indicate their dialects from a predefined list (*e.g.*, the United States English, Australian English, the United Kingdom English) or self-declare. We recommend letting participants provide their dialects in a free-form text box during data collection. For instance, a participant may state their spoken language as Arabic while they also provide their dialect as Levantine Arabic.

**Most Commonly Used Languages.** For our dataset, we took a holistic view and while we took into consideration the statistics of the most commonly used languages. We also reviewed various sources that demonstrate a compilation of the frequently used languages including Duolingo<sup>3</sup>, Web languages, Wikipedia articles (Table 4). We first checked the most spoken languages world-wide (most Top-20 spoken languages) [43] as well as nation-wide in the U.S. (Top-30 Languages

<sup>3</sup>An online language learning platform; It is the world's most-downloaded education app.; <https://www.duolingo.com>

Most commonly used languages	World most Top-20 spoken languages [43]	Top-30 spoken languages at home in US* [27]	World’s most used Top-10 languages in the Web [128]	Top-30 languages sorted by number of Wikipedia articles [45]	UN official languages [102]	Top-10 languages studied around the world from Duolingo [23]
English	✓		✓	✓	✓	✓
Spanish (Castilian and Neutral)	✓	✓	✓	✓	✓	✓
Japanese	✓	✓	✓	✓		✓
Mandarin/Cantonese/Wu (Shanghainese)	✓	✓	✓	✓	✓	✓
Korean		✓		✓		✓
Vietnamese	✓	✓		✓		
Hindi, Tamil, Telugu	✓	✓				
Arabic	✓	✓	✓	✓	✓	
Farsi						
French	✓	✓	✓	✓	✓	✓
German	✓	✓	✓	✓		✓
Italian		✓		✓		✓
Slavic languages	✓	✓	✓	✓	✓	✓
Portuguese	✓	✓	✓	✓		✓
Turkish	✓			✓		
Swahili						
Thai		✓				
Bahasa (Indonesian)	✓		✓	✓		
American Sign Language						

Table 4. Most commonly used languages based on various resources. \*This statistics excludes English.

spoken at home other than English in the U.S.) [27]. Besides the spoken languages, the most commonly used languages online are referred through statistics on World’s most used Top-10 languages in the Web (2020) [128] and Top-30 languages sorted by the number of Wikipedia articles [45]. Other references include the U.N. official languages [102] and Top-10 languages studied around the World (2021) from the Duolingo app [23]. As a result, the list provided in the Table 4 below is composed by counting the intersections of the listed sources as well as the selected languages from the research papers as follow.

**Diverse set of Languages.** We explored the diverse set of languages by referring to the literature in the field of NLP. Language understanding has always been an important task in the development of the NLP models. Since the 1990s, much research effort has been put into solving multilingual problems [81, 83, 103]. Others have provided multilingual datasets. Meta AI proposed a Multilingual LibriSpeech (MLS) dataset, which includes *eight* languages’ speech data from audio books [113]. Google Research introduced a multilingual dataset that consists of more than *forty* languages, in total around *forty billion* characters [54]. Many multilingual benchmarks, such as XGLUE [80] and XTREME [63] have been introduced to evaluate the multilingual AI models. With the development of large-scale pre-training, pre-trained multilingual models (*e.g.*, mBART [84] and mT5 [143]) show significant success in immensely complicated tasks. In addition, Meta AI has built the first unified model for the multilingual machine translation that translates directly between the *two hundred* languages without relying on the English data [47].



## 2.4 Geo-Location (Country/State/City)

As many open-source datasets are collected via crowd-sourcing, the geographical location of the speech/video data recording is not often shared. The well-known public speech datasets, Common Voice [13], M-AIILABS Speech dataset [126] and LibriSpeech [111], also do not include any geographical information but only assume the speaker of crowd sourced data is a native speaker. The Casual Conversations v1 dataset is one of the few examples that shares geo-location information where videos were recorded but it is only limited to *five* US cities.

However, geo-location can be useful information in speech and natural language data to assess the models’ robustness. For instance, as one of the contributors of linguistic variations (dialect) [21], geo-location is used to analyze AI speech problems for the same country but for different regions and/or cities. Thus, we propose instead to collect geo-location (where videos were collected) in addition to the language spoken in each video. While granularity of geo-location could be as precise as longitude/latitude, we believe country, state/region and city will be enough to assess model robustness and fairness with respect to language and dialects.

## 2.5 Disability

Building AI systems that can better serve people with disabilities has been a major aspiration within the research community [22, 139]. The key theme of prominent “Nothing about us without us” movement [31], casts light on the criticality of active inclusion of disabled participants. However, approaches for evaluating the AI models’ effectiveness for disabled people is not extensively investigated in comparison to other protected characteristics such as age and gender in audio and vision systems. [131] points out the diversity and outliers as one of the main challenges for AI fairness for people with disabilities as the disability is not *a simple variable with a small number of discrete values* unlike other attributes such as age. One individual may have several different disabilities with varying degrees, and ability may evolve over time – it requires the consideration of degree and diversity.

Since our dataset does not focus on specific disability, we considered the criteria for disability classification from different perspectives. Disability is generally classified into physical, sensory, intellectual and mental [14, 141] categories. In academic research works, there has been no clear set of categories. The category of disability itself is generally considered for assessing bias towards the people with disabilities without any subcategories labels in [100]. [124] approached with few more categories of auditory, neurological, visual and unspecific. Academic researchers have also divided disability into hierarchical categories from general to specific subgroups [52].

We propose disability categories as the following: vision (blind, low-vision, visual field loss, vision-impaired, color-blindness, loss of depth perception, glasses-wearing); hearing (deaf, hard-of-hearing, hearing with one ear, cochlear-implant-using), physical (mobility, dexterity); speech (aphasic, speech-impaired, speech loss), cognitive (Learning/Neuro-diversity); use of assistive technology.

## 2.6 Physical Adornments and Physical Attributes

Body and/or face markings are rarely taken into consideration in data collection as the collection of such data often requires more time and is a resource-intensive process. There is limited research work with physical attribute subcategories. Instead, we referred to the use cases of metaverse or online character creation such as Memoji (Apple) [12] and Avatar(Meta) [34]. We propose to divide the category into physical adornments and physical attributes as illustrated in Table 5. For each participant, we suggest asking as YES or NO for each criteria, and provide an option of free-form text for participants to provide details if they wish.



	Subcategories
Physical Adornments	Have hair cover, Hair color (dye), Have beard/mustache, Have face covering, Have face mask, Have make-up, Have eye wear (sunglasses, glasses, contact lenses, <i>etc.</i> ), Have ear wear (ear rings), Have visible tattoos, Have bindi (South Asian face markings), Have visible piercings
Physical Attributes	Hair type, Hair color (original), Hearing aids, Eye color, Have birthmarks/fleckers/moles, Have other non-tattoo facial markings

Table 5. Proposed subcategories for Physical Adornments and Attributes.

## 2.7 Voice Timbre

Voice timbre is a unique tonal quality that gives color and personality to a person’s voice. It is crucial to help make AI systems, especially Automatic Speech Recognition (ASR) models, more equally performant with regards to voice timbre. Voice timbre is generally categorized into Soprano, Mezzo Soprano, Alto, Contralto, Countertenor, Tenor, Baritone and Bass based on music theory [116, 117, 127]. Moreover, one special category for pre-pubescent voice is called Treble. Although we do not plan to include minors in our dataset, we intend to include this voice timbre label as well, as age may not be dispositive of voice timbre

In practice, voice timbre is often used as a technical term for music theory, thus, it is hard for general individuals to self-identify their own voice timbres. Academic publications either categorize voice from high-pitched to deep (*e.g.*, High-pitched, Ethereal, Sweet, Powerful, Gravelly and Deep) [122], or use gender/age-referring phrases such as middle-aged women, young men, or little girls [89]. One study utilizes a dataset that is collected from singers and separates three groups for both male and female [118] (see Table 6).

We propose three subcategories of low pitch, average pitch, and high pitch to ensure easier labeling. These subcategories provide intuitive names while enabling us to gather meaningful enough information for testing model performance in voice timbre variation.

Ours	[117]	[127]	[116]	[122]	[89]	[118]
high pitch	Soprano	Soprano	Soprano	High-pitched	Middle-aged Women	Soprano
average pitch	Mezzo	Mezzo-soprano	Mezzo Soprano	Ethereal	queen-style ladies	Mezzosoprano
low pitch	Alto	Contralto	Alto	Sweet	Young girls	Alto
	Tenor	Tenor	Contralto	Powerful	Cute little girls	Tenor
	Bass	Baritone	Countertenor	Gravelly	Elderly	Baritone
		Bass	Tenor	Deep	Middle-aged men	Bass
			Baritone		Young men	
			Bass		Cute little boys	

Table 6. Suggested voice timbre subcategories and categorization of voice timbre from music theory and academic research.

## 2.8 Apparent Skin Tone

Skin tone is an important attribute of human appearance, with significant variation from pale to dark. Recently, AI systems, especially computer vision models, have become controversial over concerns about the potential bias of

performance varying based on the skin tone [26, 72, 87]. Therefore, we include skin tone as one of the categories to collect data for the purposes of measuring algorithmic bias.

Skin tone (or skin color in some literature [17, 142]) does not have one clear universal measure even within the field of dermatology [94]. In computer vision, the Fitzpatrick scale [120] is currently the most commonly used numerical classification schema for skin tone due to its simplicity and widespread use [9, 26, 98, 140]. However, some literature points out that the use of the Fitzpatrick scale may be problematic as it is an unreliable estimator of the skin pigmentation [109]. Recently, the Monk Skin Tone (MST) Scale [96] was proposed in collaboration with Google which is going to be utilized in several Google products [65] including search and photos. In contrast to the 6-tone Fitzpatrick scale, Monk has 10 tones that favors darker and lighter skin tones equally. Many studies in different fields have used the MST Scale [50, 86, 97].

We suggest using both Fitzpatrick and Monk scales for apparent skin tone annotation. Using both scales will allow clearer comparison with previous works that use Fitzpatrick scale while also enabling measurement based on more balanced skin tone information collection through Monk scale.

### 2.9 Different Recording Setup

To understand the models’ robustness and fairness, the different recording setup is also an important element to consider as all environments in the videos cannot be uniform across different datasets – and such variation may have an underlying relationship to factors relevant to understanding unfair bias such as income level. From simple consideration of resolutions of the videos (*e.g.*, high-definition, low-definition, smartphones) to environment (*e.g.*, outdoor, indoor, beach, *etc.*), we suggest to consider varieties of recording setups as suggested in Table 7.

Recording setup	Notes
Setup	Subscripted text for speech, Unsubscripted free-form speech, Interview
Environment	Indoor, Outdoor, Beach, Field, Mountain, Office, Living room, Bedroom, Kitchen, Street, Sky, Water, <i>etc.</i>
Diverse weather settings	-
Diversity in northern and southern hemisphere (sun position)	-
Types of video quality	Web camera, Laptop camera, Smartphone camera, Professional camera, <i>etc.</i> (technical details about the camera if possible)
Background noise	-
Camera position	Hand-held phone, Selfie, Laptop, Photography studio, <i>etc.</i>
Object bounding boxes and/or segmentation	-
Lighting	Perceptually dark or light

Table 7. Different recording setup.

## 2.10 Activity

Types of human activities are broadly categorized as gestures, actions and interactions [2]. The subcategories for each activity category can be fine-grained to 200+ activities [28]. Since annotating a whole list of human activities can be innumerable, we mainly focused on the most overlapping list of activities and also the most applicable actions, gestures and appearance in our data collection setup shown in Table 8. We mainly referred to computer vision research paper that focuses on daily human activity detection [2, 11, 24, 73, 146].

Activity	Subcategories
Action	Standing
	Walking
	Sitting
	Laying
	Waving
Gesture	Stretching body
	Raising hand/leg
	Moving head
Appearance	Full body visible
	Upper body visible
	Lower body visible
	Only head visible

Table 8. List of activities. We group activities into *action*, *gesture* and *appearance*.

## 3 DELIBERATE EXCLUSION OF CERTAIN CATEGORIES

In the previous Section 2, we went over each of the ten different categories and the related literature and rationale for final sub categories. However, there are some categories that cannot be generalized to the world population and may not be relevant to the aforementioned measurement or ML training tasks. In this section, we explain the reason why we will not include them in our dataset.

**Race & Ethnicity.** Both race and ethnicity are categories that are often discussed in fairness analysis of AI systems [32, 69, 134, 138]. For instance, most studies in face recognition fairness have heavily relied on race/ethnicity [37, 69, 134]. However, Roth [119] argues that there is no single dimension that identifies the person’s “true” or “correct” race because race is being experienced as a result of conflicting dimensions including self-identification, perception by others in the society, person’s belief on others’ perception, racial self-classification, skin color, racial appearance and racial ancestry. By decoupling race into its diverse dimensions, Roth demonstrates that race is a social construct and discusses further to show racial fluidity and racial boundary change. Furthermore, race has ambiguities to be well-defined and generalized to all regions (i.e., race and ethnicity are defined differently in different countries). Barjubani *et al.* [18] also highlights that “*agreeing on a catalogue of human races has so far proved impossible.*”

Recently, a shift towards using skin tone instead is starting to happen [26, 35, 72, 135], mostly due to the problems outlined here. Moreover, when evaluating visual or auditory models it is more accurate to consider actual visual attributes [46] rather than race with its limitations and ambiguities. We already have location, apparent skin-tone and language/accent/dialect information which provide a robust baseline analysis of AI models performance. Thus, the

additional collection of controversial/sensitive data about race/ethnicity fails to override the reasons not to.

**Facial Expression.** There has been a substantial amount of research conducted in the field of *apparent emotion detection* [3, 4, 10, 15, 20, 68, 71, 76, 78, 79, 85, 88, 95, 123, 125, 145]. Nevertheless, the most recent studies have unearthed several risks of mis-contextualizing *facial expressions*. Barrett *et al.* [19] discusses the challenges of inferring *emotions* as expressions (anger, fear, disgust, happiness, *etc.*) may be communicated in many different forms across different cultures and situations. As a result, mis-contextualizing of emotions may result in potential harms in the commercial applications. Hernandez *et al.* [59] provides *twelve* guidelines for systematic assessment and reducing the risk presented by apparent emotion detection applications. Moreover, *facial expression* could be potentially utilized to reveal unintended information of individuals [58]. Considering all these risks and ambiguity in *facial expression detection*, we decided to exclude this category from the list of labels we plan to collect for the dataset.

#### 4 CONCLUSION

In this article, we presented a literature survey we have conducted for Casual Conversations v2, *i.e.* a large consent driven dataset designed for measuring algorithmic bias and robustness of AI systems. We proposed to collect *six* self-provided (by the participants) and *four* annotated categories, including *age, gender, language/dialect, geo-location (country/state/city), disability, physical adornment, physical attributes, voice timbre, apparent skin tone, different recording setup* and *activity*. For each of these categories, we provided a thorough review and discussed the proposed subcategories that we aim to collect. Furthermore, in the spirit of building a responsible dataset, three sensitive categories, *i.e. race & ethnicity* and *facial expressions*, have been excluded for ethical and safety reasons that also aligns with our ongoing efforts and long term vision in Responsible AI.

This literature review has been carried out by Meta as part of our continued commitment to Responsible AI in collaboration with academic partners. *Casual Conversations v2* will be unique in the academic world in terms of its proposed categories (which are rarely seen in the publicly available audio/vision benchmarks) and also the countries where the data collection will take place. Our dataset will enable the research community 1) to develop new models that are more fair, inclusive, and robust and also 2) to measure their models on these axes.

#### ACKNOWLEDGMENTS

We would like to express our deepest gratitude to many partners involved in the process of this literature review and also in the construction of this dataset, including but not limited to Civil Rights, Accessibility, Responsible AI, AI Analytics, Assistant, Speech Recognition, and FAIR teams at Meta. We also want to thank Eric Smith, Skyler Wang, Tashrima Hossain, Sarah Smurthwaite, Matt Bonna and Ida Cheng for their support and contribution.

#### REFERENCES

- [1] Olatunbosun Agbo-Ajala and Serestina Viriri. 2021. Deep learning approach for facial age classification: a survey of the state-of-the-art. *Artificial Intelligence Review* 54, 1 (2021), 179–213.
- [2] Jake K Aggarwal and Michael S Ryou. 2011. Human activity analysis: A review. *Acm Computing Surveys (Csur)* 43, 3 (2011), 1–43.
- [3] Niki Aifanti, Christos Papachristou, and Anastasios Delopoulos. 2010. The MUG facial expression database. In *11th International Workshop on Image Analysis for Multimedia Interactive Services WIAMIS 10*. 1–4.
- [4] Vitor Albiero, Olga RP Bellon, and Luciano Silva. 2018. Multi-label action unit detection on multiple head poses with dynamic region learning. In *2018 25th IEEE International Conference on Image Processing (ICIP)*. IEEE, 2037–2041.
- [5] Vitor Albiero, Kevin Bowyer, Kushal Vangara, and Michael King. 2020. Does face recognition accuracy get better with age? deep face matchers say no. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 261–269.

- [6] Vitor Albiero, Krishnapriya KS, Kushal Vangara, Kai Zhang, Michael C King, and Kevin W Bowyer. 2020. Analysis of gender inequality in face recognition accuracy. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*. 81–89.
- [7] Vitor Albiero, Kai Zhang, Michael C King, and Kevin W Bowyer. 2021. Gendered differences in face recognition accuracy explained by hairstyles, makeup, and facial morphology. *IEEE Transactions on Information Forensics and Security* 17 (2021), 127–137.
- [8] Alexander Amini, Ava P Soleimany, Wilko Schwarting, Sangeeta N Bhatia, and Daniela Rus. 2019. Uncovering and mitigating algorithmic bias through learned latent structure. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. 289–295.
- [9] Alexander Amini, Ava P. Soleimany, Wilko Schwarting, Sangeeta N. Bhatia, and Daniela Rus. 2019. Uncovering and Mitigating Algorithmic Bias through Learned Latent Structure. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (Honolulu, HI, USA) (AIES '19)*. Association for Computing Machinery, New York, NY, USA, 289–295. <https://doi.org/10.1145/3306618.3314243>
- [10] Deepali Aneja, Bindita Chaudhuri, Alex Colburn, Gary Faigin, Linda Shapiro, and Barbara Mones. 2018. Learning to Generate 3D Stylized Character Expressions from Humans. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 160–169. <https://doi.org/10.1109/WACV.2018.00024>
- [11] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge L Reyes-Ortiz. 2012. Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine. In *International workshop on ambient assisted living*. Springer, 216–223.
- [12] Apple. 2022. Use memojoi on your iphone or iPad pro. <https://support.apple.com/en-us/HT208986>
- [13] Rosana Ardila, Megan Branson, Kelly Davis, Michael Kohler, Josh Meyer, Michael Henretty, Reuben Morais, Lindsay Saunders, Francis Tyers, and Gregor Weber. 2020. Common Voice: A Massively-Multilingual Speech Corpus. In *Proceedings of the 12th Language Resources and Evaluation Conference*. 4218–4222.
- [14] Aruma. 2022. Types of disabilities. <https://www.aruma.com.au/about-us/about-disability/types-of-disabilities/>. Accessed: 2022-11-07.
- [15] AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Multimodal Language Analysis in the Wild: CMU-MOSEI Dataset and Interpretable Dynamic Fusion Graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Melbourne, Australia, 2236–2246. <https://doi.org/10.18653/v1/P18-1208>
- [16] World Bank. 2021. Age distribution in the United States from 2010 to 2020. <https://www.statista.com/statistics/270000/age-distribution-in-the-united-states/>
- [17] Yair Bar-Haim, Talia Saidel, and Galit Yovel. 2009. The role of skin colour in face recognition. *Perception* 38, 1 (2009), 145–148.
- [18] Guido Barbujani and Vincenza Colonna. 2011. *Genetic Basis of Human Biodiversity: An Update*. Springer Berlin Heidelberg, Berlin, Heidelberg, 97–119. [https://doi.org/10.1007/978-3-642-20992-5\\_6](https://doi.org/10.1007/978-3-642-20992-5_6)
- [19] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M. Martinez, and Seth D. Pollak. 2019. Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements. *Psychological Science in the Public Interest* 20, 1 (2019), 1–68. <https://doi.org/10.1177/1529100619832930> arXiv:<https://doi.org/10.1177/1529100619832930> PMID: 31313636.
- [20] Júlio César Batista, Vitor Albiero, Olga RP Bellon, and Luciano Silva. 2017. Aumpnet: simultaneous action units detection and intensity estimation on multipose facial images using a single convolutional neural network. In *2017 12th IEEE international conference on automatic face & gesture recognition (FG 2017)*. IEEE, 866–871.
- [21] Joan C Beal. 2010. *Introduction to Regional Englishes*. Edinburgh University Press.
- [22] Cynthia L Bennett and Os Keyes. 2020. What is the point of fairness? Disability, AI and the complexity of justice. *ACM SIGACCESS Accessibility and Computing* 125 (2020), 1–1.
- [23] Cindy Blanco. 2021. The 2021 Duolingo Language Report. <https://blog.duolingo.com/2021-duolingo-language-report/>
- [24] Aaron F. Bobick and James W. Davis. 2001. The recognition of human movement using temporal templates. *IEEE Transactions on pattern analysis and machine intelligence* 23, 3 (2001), 257–267.
- [25] Janet Brito. 2018. Genderqueer: What does it mean? <https://www.healthline.com/health/transgender/genderqueer>
- [26] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [27] United States Census Bureau. 2017. Top-30 Languages spoken at home other than English in US. <https://data.census.gov/cedsci/>. Accessed: 2022-11-07.
- [28] Fabian Caba Heilbron, Victor Escorcía, Bernard Ghanem, and Juan Carlos Niebles. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 961–970.
- [29] Statistics Canada. 2017. Age Categories, Life Cycle Groupings. <https://www.statcan.gc.ca/en/concepts/definitions/age2>
- [30] LGBTQIA Resource Center. 2022. LGBTQIA Resource Center Glossary. <https://lgbtqia.ucdavis.edu/educated/glossary>. Accessed: 2022-11-07.
- [31] James I. Charlton. 1998. *Nothing About Us Without Us: Disability Oppression and Empowerment* (1 ed.). University of California Press. <http://www.jstor.org/stable/10.1525/j.ctt1pnqn9>
- [32] Xingyu Chen, Zhengxiong Li, Srirangaraj Setlur, and Wenyao Xu. 2022. Exploring racial and gender disparities in voice biometrics. *Scientific Reports* (Mar 2022). <https://doi.org/10.1038/s41598-022-06673-y>
- [33] United States Equal Employment Opportunity Commission. 2022. EEOC to add non-binary gender option to discrimination charge intake process. <https://www.eeoc.gov/newsroom/eeoc-add-non-binary-gender-option-discrimination-charge-intake-process>. Accessed: 2022-11-07.
- [34] Josh Constine. 2019. Facebook introduces Avatars, its Bitmoji competitor. <https://techcrunch.com/2019/06/03/facebook-avatars-stickers/>

- [35] Cynthia M. Cook, John J. Howard, Yevgeniy B. Sirotnin, Jerry L. Tipton, and Arun R. Vemury. 2019. Demographic Effects in Facial Recognition and Their Dependence on Image Acquisition: An Evaluation of Eleven Commercial Systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science* (2019).
- [36] Kate Crawford and Trevor Paglen. 2021. Excavating AI: the politics of images in machine learning training sets. *AI & SOCIETY* (Dec 2021). <https://doi.org/10.1007/s00146-021-01162-8>
- [37] Abhijit Das, Antitza Dantcheva, and Francois Bremond. 2018. Mitigating bias in gender, age and ethnicity classification: a multi-task convolution neural network approach. In *Proceedings of the european conference on computer vision (eccv) workshops*. 0–0.
- [38] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [39] Yuan Dong, Yanan Liu, and Shiguo Lian. 2016. Automatic age estimation based on deep learning algorithm. *Neurocomputing* 187 (2016), 4–10.
- [40] Samuel Dooley, George Zhihong Wei, Tom Goldstein, and John P Dickerson. 2022. Robustness Disparities in Face Detection. In *Neural Information Processing Systems Datasets and Benchmarks Track*. <https://openreview.net/forum?id=UXPXs-OYbks>
- [41] Mengnan Du, Subhabrata Mukherjee, Guanchu Wang, Ruixiang Tang, Ahmed Awadallah, and Xia Hu. 2021. Fairness via representation neutralization. *Advances in Neural Information Processing Systems* 34 (2021), 12091–12103.
- [42] Mingxing Duan, Kenli Li, Canqun Yang, and Keqin Li. 2018. A hybrid deep learning CNN-ELM for age and gender classification. *Neurocomputing* 275 (2018), 448–461.
- [43] David M. Eberhard, Gary F. Simons, and Charles D. Fenning. 2022. What are the top 200 most spoken languages? <https://www.ethnologue.com/guides/ethnologue200>
- [44] The Economist. 2022. Which countries offer gender-neutral passports? <https://www.economist.com/graphic-detail/2022/04/11/which-countries-offer-gender-neutral-passports>. Accessed: 2022-11-07.
- [45] Wikipedia: The Free Encyclopedia. 2022. List of wikipedias. [https://en.wikipedia.org/wiki/List\\_of\\_Wikipedias](https://en.wikipedia.org/wiki/List_of_Wikipedias). Accessed: 2022-11-07.
- [46] Simone Fabbrizzi, Symeon Papadopoulos, Eirini Ntoutsi, and Ioannis Kompatsiaris. 2022. A survey on bias in visual datasets. *Computer Vision and Image Understanding* 223 (2022), 103552. <https://doi.org/10.1016/j.cviu.2022.103552>
- [47] Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. 2021. Beyond English-Centric Multilingual Machine Translation. *J. Mach. Learn. Res.* 22, 107 (2021), 1–48.
- [48] Siyuan Feng, Olya Kudina, Bence Mark Halpern, and Odette Scharenborg. 2021. Quantifying bias in automatic speech recognition. *arXiv preprint arXiv:2103.15122* (2021).
- [49] World Economic Forum. 2022. Artificial Intelligence for Children. [https://www3.weforum.org/docs/WEF\\_Artificial\\_Intelligence\\_for\\_Children\\_2022.pdf](https://www3.weforum.org/docs/WEF_Artificial_Intelligence_for_Children_2022.pdf).
- [50] Steven L Foy, Victor Ray, and Ashley Hummel. 2017. The shade of a criminal record: Colorism, incarceration, and external racial classification. *Socius* 3 (2017), 2378023116689567.
- [51] Priya Goyal, Adriana Romero Soriano, Caner Hazirbas, Levent Sagun, and Nicolas Usunier. 2022. Fairness Indicators for Systematic Assessments of Visual Feature Extractors. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. <https://doi.org/10.1145/3531146.3533074>
- [52] Kristin M Graham, Brian T McMahon, Jeong Han Kim, Paige Simpson, and Megan C McMahon. 2019. Patterns of workplace discrimination across broad categories of disability. *Rehabilitation psychology* 64, 2 (2019), 194.
- [53] Patrick Grother, Mei Ngan, and Kayee Hanaoka. 2018. Ongoing Face Recognition Vendor Test (FRVT) Part 2: Identification.
- [54] Mandy Guo, Zihang Dai, Denny Vrandečić, and Rami Al-Rfou. 2020. Wiki-40b: Multilingual language model dataset. In *Proceedings of the 12th Language Resources and Evaluation Conference*. 2440–2452.
- [55] Xudong Han, Timothy Baldwin, and Trevor Cohn. 2022. Towards Equal Opportunity Fairness through Adversarial Learning. *arXiv preprint arXiv:2203.06317* (2022).
- [56] Margot Hanley, Apoorv Khandelwal, Hadar Averbuch-Elor, Noah Snively, and Helen Nissenbaum. 2020. An ethical highlighter for people-centric dataset creation. *arXiv preprint arXiv:2011.13583* (2020).
- [57] Caner Hazirbas, Joanna Bitton, Brian Dolhansky, Jacqueline Pan, Albert Gordo, and Cristian Canton Ferrer. 2022. Towards Measuring Fairness in AI: The Casual Conversations Dataset. *IEEE Transactions on Biometrics, Behavior, and Identity Science* (2022).
- [58] Stig Hebbelstrup, Rye Rasmussen, Steven ludeke, and Robert Klemmensen. 2022. Using deep learning to predict ideology from facial photographs: Expressions, beauty, and extra-facial information. (2022). <https://doi.org/10.31234/osf.io/5mjuq> psyarxiv:https://psyarxiv.com/5mjuq/
- [59] Javier Hernandez, Josh Lovejoy, Daniel McDuff, Jina Suh, Tim Brien, Arathi Sethumadhavan, Gretchen Greene, Rosalind Picard, and Mary Czerwinski. 2021. Guidelines for Assessing and Minimizing Risks of Emotion Recognition Applications. In *International Conference on Affective Computing & Intelligent Interaction (ACII 2021)*. <https://www.microsoft.com/en-us/research/publication/guidelines-for-assessing-and-minimizing-risks-of-emotion-recognition-applications/>
- [60] Wen-Bing Horng, Cheng-Ping Lee, Chun-Wen Chen, et al. 2001. Classification of age groups based on facial features. *Journal of Applied Science and Engineering* 4, 3 (2001), 183–192.
- [61] The White House. 2022. Readout of White House Listening Session on Tech Platform Accountability. <https://www.whitehouse.gov/briefing-room/statements-releases/2022/09/08/readout-of-white-house-listening-session-on-tech-platform-accountability/>. Accessed: 2022-09-19.
- [62] John J. Howard, Yevgeniy B. Sirotnin, and Arun R. Vemury. 2019. The Effect of Broad and Specific Demographic Homogeneity on the Imposter Distributions and False Match Rates in Face Recognition Algorithm Performance. In *International Conference on Biometrics Theory, Applications and*



## Systems.

- [63] Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*. PMLR, 4411–4421.
- [64] Jin Huang, Bin Li, Jia Zhu, and Jian Chen. 2017. Age classification with deep learning face representation. *Multimedia Tools and Applications* 76, 19 (2017), 20231–20247.
- [65] Alphabet Inc. 2022. A closer look at the research to help AI see more skin tones. <https://blog.google/technology/research/ai-monk-scale-skin-tone-story/>. Accessed: 2022-11-03.
- [66] Mark Jordan. 2021. How to write gender questions for a survey. <https://www.alchemer.com/resources/blog/how-to-write-survey-gender-questions/>
- [67] Sangwon Jung, Sanghyuk Chun, and Taesup Moon. 2022. Learning Fair Classifiers with Partially Annotated Group Labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10348–10357.
- [68] Miyuki Kamachi, Michael Lyons, and Jiro Gyoba. 1997. The japanese female facial expression (jaffe) database. Available: <http://www.kasrl.org/jaffe.html> (01 1997).
- [69] Kimmo Karkkainen and Jungseock Joo. 2021. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1548–1558.
- [70] Mehtab Khan and Alex Hanna. 2023. The Subjects and Stages of AI Dataset Development: A Framework for Dataset Accountability. *Forthcoming 19 Ohio St. Tech. L.J.* (2023). <https://ssrn.com/abstract=4217148>
- [71] Dimitrios Kollias and Stefanos Zafeiriou. 2019. Expression, Affect, Action Unit Recognition: Aff-Wild2, Multi-Task Learning and ArcFace. *CoRR* abs/1910.04855 (2019). arXiv:1910.04855 <http://arxiv.org/abs/1910.04855>
- [72] KS Krishnapriya, Vitor Albiero, Kushal Vangara, Michael C King, and Kevin W Bowyer. 2020. Issues related to face recognition accuracy varying based on race and skin tone. *IEEE Transactions on Technology and Society* 1, 1 (2020), 8–20.
- [73] Ku Nurhanim Ku Abd. Rahim, I Elamvazuthi, Lila Iznita Izhar, and Genci Capi. 2018. Classification of human daily activities using ensemble methods based on smartphone inertial sensors. *Sensors* 18, 12 (2018), 4132.
- [74] Neeraj Kumar, Alexander C Berg, Peter N Belhumeur, and Shree K Nayar. 2009. Attribute and simile classifiers for face verification. In *2009 IEEE 12th international conference on computer vision*. IEEE, 365–372.
- [75] Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhi Wang, and Ed Chi. 2020. Fairness without demographics through adversarially reweighted learning. *Advances in neural information processing systems* 33 (2020), 728–740.
- [76] Oliver Langner, Ron Dotsch, Gijsbert Bijlstra, Daniel H. J. Wigboldus, Skyler T. Hawk, and Ad van Knippenberg. 2010. Presentation and validation of the Radboud Faces Database. *Cognition and Emotion* 24, 8 (2010), 1377–1388. <https://doi.org/10.1080/02699930903485076> arXiv:<https://doi.org/10.1080/02699930903485076>
- [77] Nayeon Lee, Yejin Bang, Jamin Shin, and Pascale Fung. 2019. Understanding the Shades of Sexism in Popular TV Series. In *Proceedings of the 2019 Workshop on Widening NLP*. 122–125.
- [78] Shan Li and Weihong Deng. 2019. Reliable Crowdsourcing and Deep Locality-Preserving Learning for Unconstrained Facial Expression Recognition. *IEEE Transactions on Image Processing* 28, 1 (2019), 356–370.
- [79] Shan Li, Weihong Deng, and JunPing Du. 2017. Reliable Crowdsourcing and Deep Locality-Preserving Learning for Expression Recognition in the Wild. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2584–2593.
- [80] Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, et al. 2020. Xglue: A new benchmark dataset for cross-lingual pre-training, understanding and generation. *arXiv preprint arXiv:2004.01401* (2020).
- [81] Marina Litvak, Mark Last, and Menahem Friedman. 2010. A new approach to improving multilingual summarization using a genetic algorithm. In *Proceedings of the 48th annual meeting of the association for computational linguistics*. 927–936.
- [82] Chunxi Liu, Michael Picheny, Leda Sari, Pooja Chitkara, Alex Xiao, Xiaohui Zhang, Mark Chou, Andres Alvarado, Caner Hazirbas, and Yatharth Saraf. 2022. Towards Measuring Fairness in Speech Recognition: Casual Conversations Dataset Transcriptions. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [83] Jingjing Liu, Panupong Pasupat, Scott Cyphers, and Jim Glass. 2013. Asgard: A portable architecture for multilingual dialogue systems. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 8386–8390.
- [84] Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics* 8 (2020), 726–742.
- [85] Steven R. Livingstone and Frank A. Russo. 2018. The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): A dynamic, multimodal set of facial and vocal expressions in North American English. *PLOS ONE* 13, 5 (05 2018), 1–35. <https://doi.org/10.1371/journal.pone.0196391>
- [86] Patricia Louie and Rima Wilkes. 2018. Representations of race and skin tone in medical textbook imagery. *Social Science & Medicine* 202 (2018), 38–42.
- [87] Boyu Lu, Jun-Cheng Chen, Carlos D Castillo, and Rama Chellappa. 2019. An experimental evaluation of covariates effects on unconstrained face verification. *IEEE Transactions on Biometrics, Behavior, and Identity Science* 1, 1 (2019), 42–55.
- [88] Patrick Lucey, Jeffrey F. Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. 2010. The Extended Cohn-Kanade Dataset (CK+): A complete dataset for action unit and emotion-specified expression. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*. 94–101. <https://doi.org/10.1109/CVPRW.2010.5543262>



- [89] Yinyu Lyu, Chenwei Liu, Hao Tan, Rui Xie, Keke Tang, and Zhaoquan Gu. 2020. Convolutional Neural Network based Timbre Classification. In *Proceedings of the 2020 International Conference on Cyberspace Innovation of Advanced Technologies*. 186–191.
- [90] Xinyao Ma, Maarten Sap, Hannah Rashkin, and Yejin Choi. 2020. PowerTransformer: Unsupervised Controllable Revision for Biased Language Correction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 7426–7441.
- [91] California State University San Marcos. 2022. Inclusive language guidelines: Gender identity. <https://www.csusm.edu/ipa/surveys/inclusive-language-guidelines.html>. Accessed: 2022-11-07.
- [92] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.* (2021). <https://doi.org/10.1145/3457607>
- [93] Zion Mengesha, Courtney Heldreth, Michal Lahav, Juliana Sublewski, and Elyse Tuennerman. 2021. “I don’t Think These Devices are Very Culturally Sensitive.”—Impact of Automated Speech Recognition Errors on African Americans. *Frontiers in Artificial Intelligence* (2021). <https://www.frontiersin.org/articles/10.3389/frai.2021.725911>
- [94] Michele Merler, Nalini Ratha, Rogerio S Feris, and John R Smith. 2019. Diversity in faces. *arXiv preprint arXiv:1901.10436* (2019).
- [95] Ali Mollahosseini, Behzad Hassani, and Mohammad H. Mahoor. 2017. AffectNet: A Database for Facial Expression, Valence, and Arousal Computing in the Wild. *CoRR* abs/1708.03985 (2017). arXiv:1708.03985 <http://arxiv.org/abs/1708.03985>
- [96] Ellis P Monk Jr. 2014. Skin tone stratification among Black Americans, 2001–2003. *Social Forces* 92, 4 (2014), 1313–1337.
- [97] Ellis P Monk Jr. 2015. The cost of color: Skin color, discrimination, and health among African-Americans. *Amer. J. Sociology* 121, 2 (2015), 396–444.
- [98] Vidya Muthukumar, Tejaswini Pedapati, Nalini Ratha, Prasanna Sattigeri, Chai-Wah Wu, Brian Kingsbury, Abhishek Kumar, Samuel Thomas, Aleksandra Mojsilovic, and Kush R Varshney. 2018. Understanding Unequal Gender Classification Accuracy from Face Images. *arXiv e-prints* (2018), arXiv–1812.
- [99] Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. StereoSet: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. 5356–5371.
- [100] Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. CrowS-Pairs: A Challenge Dataset for Measuring Social Biases in Masked Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online.
- [101] United Nations. 1982. Provisional Guidelines on Standard International Age Classifications. [https://unstats.un.org/unsd/publication/SeriesM/SeriesM\\_74e.pdf](https://unstats.un.org/unsd/publication/SeriesM/SeriesM_74e.pdf). Accessed: 2022-11-07.
- [102] United Nations. 2022. Official Languages. <https://www.un.org/en/our-work/official-languages>. Accessed: 2022-11-07.
- [103] Joel Nothman, Nicky Ringland, Will Radford, Tara Murphy, and James R Curran. 2013. Learning multilingual named entity recognition from Wikipedia. *Artificial Intelligence* 194 (2013), 151–175.
- [104] United States Department of State — Bureau of Consular Affairs. 2022. Selecting your gender marker - travel. <https://travel.state.gov/content/travel/en/passports/need-passport/selecting-your-gender-marker.html>. Accessed: 2022-11-07.
- [105] Australian Bureau of Statistics. 2021. Standard for sex, gender, variations of sex characteristics and sexual orientation variables. <https://www.abs.gov.au/statistics/standards/standard-sex-gender-variations-sex-characteristics-and-sexual-orientation-variables/latest-release>. Accessed: 2022-11-07.
- [106] Singapore Department of Statistics. 2020. National Statistical Standards. [https://www.singstat.gov.sg/-/media/files/standards\\_and\\_classifications/nsa.pdf](https://www.singstat.gov.sg/-/media/files/standards_and_classifications/nsa.pdf). Accessed: 2022-11-07.
- [107] The University of Utah. 2019. Gender Inclusive Language in Research. [https://irb.utah.edu/\\_resources/documents/pdf/IGS%20-%20Gender%20Inclusive%20Language%20in%20Research%20081519.pdf](https://irb.utah.edu/_resources/documents/pdf/IGS%20-%20Gender%20Inclusive%20Language%20in%20Research%20081519.pdf) Accessed: 2022-11-07.
- [108] UK Government Equalities Office. 2019. National LGBT survey: Summary report. <https://www.gov.uk/government/publications/national-lgbt-survey-summary-report/national-lgbt-survey-summary-report>. Accessed: 2022-11-07.
- [109] UK Okoji, SC Taylor, and JB Lipoff. 2021. Equity in skin typing: why it is time to replace the Fitzpatrick scale. *British Journal of Dermatology* 185, 1 (2021), 198–199.
- [110] World Health Organization. 2022. Gender and health. [https://www.who.int/health-topics/gender#tab=tab\\_1](https://www.who.int/health-topics/gender#tab=tab_1). Accessed: 2022-11-7.
- [111] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 5206–5210.
- [112] Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. Reducing Gender Bias in Abusive Language Detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2799–2804.
- [113] Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. MLS: A Large-Scale Multilingual Dataset for Speech Research. (2020).
- [114] A. Ragni, K.M. Knill, S.P. Rath, and M.J.F. Gales. 2014. Data augmentation for low resource languages. , 810–814 pages. <https://eprints.whiterose.ac.uk/152844/> © 2014 International Speech Communication Association (ISCA). Reproduced in accordance with the publisher’s self-archiving policy..
- [115] Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable Auditing: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products. In *AAAI/ACM Conference on AI, Ethics, and Society*.
- [116] Matt Ramsey. 2022. The 8 Singing Voice Types: Find Out Yours Here! <https://ramseyvoice.com/voice-types/>
- [117] Tracy Reina. 2018. types of vocal timbre. <https://www.musictoyourhome.com/blog/types-of-vocal-timbre/>

- [118] Friederike Roers, Dirk Mürbe, and Johan Sundberg. 2009. Voice classification and vocal tract of singers: a study of x-ray images and morphology. *The Journal of the Acoustical Society of America* 125, 1 (2009), 503–512.
- [119] Wendy D. Roth. 2016. The multiple dimensions of race. *Ethnic and Racial Studies* 39, 8 (2016), 1310–1338. <https://doi.org/10.1080/01419870.2016.1140793>
- [120] Silonie Sachdeva et al. 2009. Fitzpatrick skin typing: Applications in dermatology. *Indian journal of dermatology, venereology and leprology* 75, 1 (2009), 93.
- [121] Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2020. Social Bias Frames: Reasoning about Social and Power Implications of Language. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 5477–5490.
- [122] Cheng-Ya Sha, Yi-Hsuan Yang, Yu-Ching Lin, and Homer H Chen. 2013. Singing voice timbre classification of Chinese popular music. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 734–738.
- [123] Shivendra Singh and Shajulin Benedict. 2020. Indian Semi-Acted Facial Expression (iSAFE) Dataset for Human Emotions Recognition. In *Advances in Signal Processing and Intelligent Recognition Systems*, Sabu M. Thampi, Rajesh M. Hegde, Sri Krishnan, Jayanta Mukhopadhyay, Vipin Chaudhary, Oge Marques, Selwyn Piramuthu, and Juan M. Corchado (Eds.). Springer Singapore, Singapore, 150–162.
- [124] Eric Michael Smith, Melissa Hall Melanie Kambadur, Eleonora Presani, and Adina Williams. 2022. "I'm sorry to hear that": finding bias in language models with a holistic descriptor dataset. *arXiv preprint arXiv:2205.09209* (2022).
- [125] Ian Sneddon, Margaret McRorie, Gary McKeown, and Jennifer Hanratty. 2012. The Belfast Induced Natural Emotion Database. *IEEE Transactions on Affective Computing* 3, 1 (2012), 32–41. <https://doi.org/10.1109/T-AFFC.2011.26>
- [126] Imdat Solak. 2019. The M-AILABS speech dataset. <https://www.caito.de/2019/01/03/the-m-ailabs-speech-dataset/>
- [127] MasterClass staff. 2022. Types of Vocal Timbre: How to Improve Your Vocal Tone. <https://www.masterclass.com/articles/vocal-timbre#quiz-0>
- [128] Internet World Stats. 2020. Top Ten internet languages in the world - internet statistics. <https://www.internetworldstats.com/stats7.htm>. Accessed: 2022-11-07.
- [129] Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating Gender Bias in Natural Language Processing: Literature Review. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 1630–1640.
- [130] Huan Tian, Tianqing Zhu, Wei Liu, and Wanlei Zhou. 2022. Image Fairness in Deep Learning: Problems, Models, and Challenges. *Neural Comput. Appl.* (2022).
- [131] Shari Trewin. 2018. AI Fairness for People with Disabilities: Point of View. *CoRR* abs/1811.10670 (2018). arXiv:1811.10670 <http://arxiv.org/abs/1811.10670>
- [132] Graeme Trousdale. 2010. *Introduction to English Sociolinguistics*. Edinburgh University Press.
- [133] Peter Trudgill and Jack K Chambers. 2017. *Dialects of English: Studies in grammatical variation*. Routledge.
- [134] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. 2019. Racial Faces in the Wild: Reducing Racial Bias by Information Maximization Adaptation Network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- [135] Mei Wang, Yaobin Zhang, and Weihong Deng. 2021. Meta Balanced Network for Fair Face Recognition. *IEEE transactions on pattern analysis and machine intelligence* (2021).
- [136] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. 2019. Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE Computer Society, 5309–5318.
- [137] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. 2020. Towards Fairness in Visual Recognition: Effective Strategies for Bias Mitigation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, 8916–8925.
- [138] Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. 2020. Towards Fairness in Visual Recognition: Effective Strategies for Bias Mitigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [139] Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kaziunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. 2019. Disability, bias, and AI. *AI Now Institute* (2019).
- [140] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. 2019. Predictive inequity in object detection. *arXiv preprint arXiv:1902.11097* (2019).
- [141] Disabled World. 2004. Disabled World. <https://www.disabled-world.com>. Accessed: 2022-11-07.
- [142] Han Xu, Abhijit Sarkar, and A Lynn Abbott. 2022. Color Invariant Skin Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2906–2915.
- [143] Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2020. mT5: A massively multilingual pre-trained text-to-text transformer. *arXiv preprint arXiv:2010.11934* (2020).
- [144] Sirui Yao and Bert Huang. 2017. Beyond parity: Fairness objectives for collaborative filtering. *Advances in neural information processing systems* 30 (2017).
- [145] Guoying Zhao, Xiaohua Huang, Matti Taini, Stan Z. Li, and Matti Pietikäinen. 2011. Facial expression recognition from near-infrared videos. *Image and Vision Computing* 29, 9 (2011), 607–619. <https://doi.org/10.1016/j.imavis.2011.07.002>

- [146] Yu Zhao, Renhong Yang, Guillaume Chevalier, and Maoguo Gong. 2017. Deep Residual Bidir-LSTM for Human Activity Recognition Using Wearable Sensors. *CoRR* abs/1708.08989 (2017). arXiv:1708.08989 <http://arxiv.org/abs/1708.08989>