"#*DisabledOnIndianTwitter*" : A Dataset towards Understanding the Expression of People with Disabilities on Indian Twitter

♤Ishani Mondal,* ♡Sukhnidh Kaur, ♡Kalika Bali, ♢Aditya Vashistha, ♡Manohar Swaminathan
♤University of Maryland, College Park, Maryland, USA
♡Microsoft Research Labs, Bangalore, India
♢Cornell University, Ithaca, New York, USA
♤imonda@umd.edu, {t-sukkaur, kalikab, swmanoh}@microsoft.com
adityav@cornell.edu

Abstract

Twitter serves as a powerful tool for self-expression among people with disabilities. To understand how disabled Indians use Twitter, we introduce a manually annotated corpus **#DisabledOnIndianTwitter** comprising 2,384 tweets published by 27 female and 15 male users. To examine patterns in Twitter use, we propose a novel hierarchical annotation taxonomy to classify the tweets into various themes including discrimination, advocacy, and self-identification. Using these annotations, we benchmark the corpus leveraging state-of-the-art classifiers. We use a mixed-methods analysis to showcase differences in self-expression among male and female disabled users.

1 Introduction

A majority of disabled Indians exist at the margins of society with little to no access to social media (Census, 2011). Structural embeddings of ableism and patriarchy further intersect to produce multiply oppressive conditions for Indian women with disabilities (Thomas and Thomas, 2002; Dawn, 2014). However, as access to ICTs and high-speed internet grows, Indian Twitter's user base is expanding to include disability influencers, activists, and everyday disabled users. Recent work from the West examines elements of 'Disability Twitter' (Hemsley et al., 2015; Mann, 2018; Ineland et al., 2019; Ellis and Goggin, 2013) but little is known of such Twitter use among the Indian user base. To fill this gap, we study self-expression among disabled users on Indian Twitter. As researchers with overlapping interests in disability and gender, we orient our analysis towards gendered self-expression.

We introduce a novel human-annotated corpus **#DisabledOnIndianTwitter** comprising 2,384 tweets published by a preliminary set of 15 male and 27 female disabled people who are active and vocal on Indian Twitter. In order to linguistically analyze the patterns of their social media usage, we propose a hierarchical linguistic annotation framework which takes into account contextual nuances surrounding disability-related concerns. Within this framework, we propose multi-level and multi-class thematic classifications including discrimination, advocacy, harassment, and self-identification. As a next step, we benchmark state-of-the-art language model classifiers fine-tuned on these datasets, noting significant room for improvement of models. Through our mixed-methods analysis on the created corpus, we find that disabled women are more likely than disabled men to center personal experiences while expressing discrimination, advocacy and harassment, but disabled men tend to be more authoritative in their expression. Disabled women in Sports are more likely to advocate for inclusion rights than disabled women in other professions.

In sum, our work makes three major contributions to NLP and accessibility research.

1. We propose a novel hierarchical annotation taxonomy to perform linguistic analysis of disability-related textual content.

2. We introduce the first-of-its-kind human-annotated dataset aimed at understanding online expressions of disabled people in India. The dataset will also serve as a baseline for future explorations on tweets generated by a demographic severely underrepresented in current NLP advances.

3. We perform a mixed-method analysis on our corpus to identify gendered differences in self-expression of disabled people on Indian Twitter.

2 Background and Related Work

<u>**Social Media and Disability:**</u> Disabled people remain one of the most disenfranchised demograph-

*This work was done when the author was a Research Fellow at Microsoft Research India.

ics across low-income countries including India (Groce et al., 2011; Buettgen et al., 2015; Pinilla-Roncancio et al., 2020). Abject poverty coupled with complex sociocultural norms pose significant barriers to equal rights and representation (Kapur Mehta and Shah, 2003; Groce et al., 2011). Disabled women in India experience vagaries of marginalization, often in the form of gender and ableist discrimination, poverty, and inadequate family support, among others (Dhungana, 2006). For example, Leveille et al. (2000) found a substantial gender differential in self-reported health where disabled women reported poorer health than their male counterparts. Scholars have also studied the struggles of people with disabilities in accessing basic education (Croft, 2013; Jameel, 2011), employment (Dyaram and V., 2020; Kumar et al., 2012), and healthcare services (Mactaggart et al., 2016).

With the rising influence of social media on everyday lives, several scholars have started studying the use and non-use of social media platforms by people with disabilities (Outini, 2020; Vashistha et al., 2015). For example, scholars have examined the impact of social media on agency and representation of people with vision impairments (Pal et al., 2017) as well as its use during public health crises (Mont et al., 2021; Mehrotra, 2021). However, there is a scarcity of research on gender differences in how disabled people engage with others on social media.

NLP Methods and Datasets on Disability: NLP Researchers have recently been focusing on studying unintended biases in NLP models against several historically marginalized groups such as those based on differences in race, culture,and gender (Bolukbasi et al., 2016; Jentzsch et al., 2019; Garg et al., 2019; Barocas et al., 2018; Dixon et al., 2018). Several datasets have been created with the goal of fostering research in quantifying societal bias, i.e., the under-representation of these demographics in NLP models that can be detrimental for downstream NLP tasks (Levy et al., 2021; Babaeianjelodar et al., 2020; Zhao et al., 2020; Sharma et al., 2021; Dinan et al., 2019). Although some recent work (Hutchinson et al., 2020; Hassan et al., 2021) has focused on quantifying the representation of the people with disabilities in pre-trained language models, there has been a general lack of attention towards building datasets to understand how disabled people engage and express themselves on social media. (Mack et al., 2021). Our work fills this gap by: (1) creating a new dataset containing tweets from people with disabilities in India and making it publicly available[1], and (2) analyzing the dataset to identify differential patterns of Twitter usage based on gender, and other attributes.

3 Dataset

We used Twitter to collect public data since it allows such analysis through APIs available for researchers. We manually selected Twitter accounts where users disclosed their disability identity, for example, in their Twitter bio, profile picture, username, display name, or within the content of their tweets. We note that disabled representation on Indian Twitter is marginal due to a lack of access to ICTs and high-speed internet among the disabled population. Further, due to the stigma associated with disability, a limited number of users disclose their disability identity on Indian Twitter. So while the manual process ensures that our dataset is accurate, it also means that we have a limited number of Twitter accounts to analyze.

Selecting Twitter Handles: We refer to our dataset as **"DisabledOnIndianTwitter"** which comprises 27 females and 15 males working as sportspersons, social workers, researchers, bloggers, actors, writers, travelers, company directors, comedians, and students. We identified occupations and genders through manually examination of Twitter bios, tweets, and other profiles.Table 5 in Appendix shows the details of these Twitter handles without explicitly disclosing their identity for privacy concerns.

Data Filtering: Next, we crawled recent tweets (last 3206 tweets per user) posted by each user. After collecting the tweets, we excluded those with duplicate or no meaningful textual content (e.g., only @-mentions or images). We only selected tweets in English using the language code provided by Twitter. During data filtration, we manually verified the language codes and excluded non-English tweets. We also excluded retweets and replies as these do not necessarily express the thoughts of the user who retweeted them. We thus obtained a set of 60,000 tweets.

[1] `https://github.com/Ishani-Mondal/-DisabledOnIndianTwitter`

Tweet	Relatedness	Discrim	Advocacy	Incl	Identity	Factual	Stance	Haras	Theme
In small industrial district of Karur: a grand beginning of accessaudits with famous Lord Murugan temple. Glad 2 see. Awed by 365 steps-solutions to make it accessible for people with disabilites	R1	D0	A0	I1	Id0	O	P	H0	O

Table 1: An Example of Annotated Tweet from our corpus. Here Discrim indicates Discrimination, Incl indicates Inclusion, Haras indicates Harassment (Shortened due to space constraint).

Categories	Statistics
Relatedness	1518 (R1), 866 (R0)
Discrimination	426 (D1), 1092 (D0)
Advocacy	638 (A1), 880 (A0)
Inclusion	186 (I1), 1332 (I0)
Identity	363 (Id1), 1155 (Id0)
Fact/Opinion	370 (F), 1148 (O)
Stance	664 (P), 484 (N)
Harassment	148 (H1), 1370 (H0)
Theme	198 (HH), 45 (Emp), 85 (Ed), 1190 (O)

Table 2: Final Statistics of our Dataset

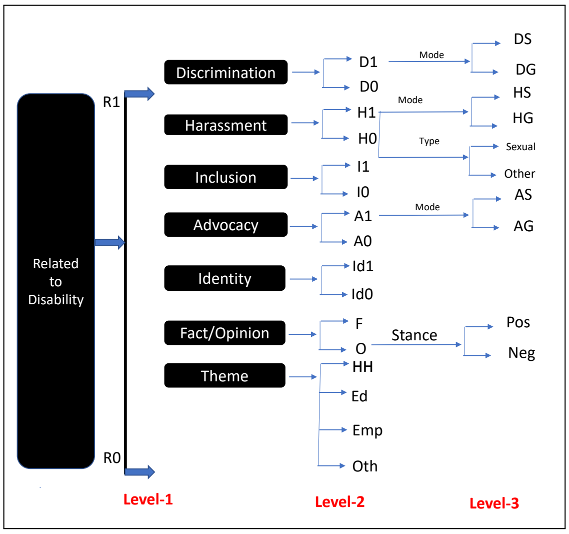

Figure 1: Hierarchical Annotation Taxonomy (Three Levels) used to thematically classify the Tweets posted by disabled people on Indian Twitter.

Keyword Based Sampling: We used a keyword-based sampling method to increase the hit rate of tweets with disability related concerns, following the existing work on labeling infrequent linguistic phenomena, e.g., irony (Van Hee et al., 2018) , hate speech (Waseem and Hovy, 2016) or bragging (Jin et al., 2022). To ensure that we capture all disability related information in the posts, we extended the list of disability related keywords provided by (Hutchinson et al., 2020) and their synonyms from WordNet (Miller, 1995). The complete list of keywords is available in Appendix 12. We observed that some tweets did not explicitly contain the keywords, but frequently mentioned accessing education, societal aspects of livelihood such as employment, e.g. *'job*', 'employ*', 'government'*, and health and hygiene, e.g. *university, education, studies*. We have selected the Tweets based on these words and added the keywords to the list shown in Appendix 12.

4 Annotation Taxonomy

We propose a linguistic annotation schema (Figure 1) to study the patterns of self-expression of tweets. The purpose is to categorize each tweet into different classes with each category indicating one of the aspects mentioned above. In this section, we define the broader and fine-grained sub-categories under each category.

4.1 Relatedness

We began the annotation exercise by determining whether each tweet contained disability-relevant subject matter. Annotators were asked to mark related tweets as (R1) and unrelated tweets as (R0). **Example (annotation R1)**: "*Its always amusing when people feel unsettled when they are around a disabled person. They just do not know what to do*

with themselves when it comes to offering support in a dignified way."
Example (annotation R0): "*Dear Pediatric Surgeons, its high time you embrace; STOP using pathologising terms like gender dysphoria, Disorders of Sex Development; Differences of Sex Dev.*"
For disability related tweets, annotators were asked to further annotate the following aspects:

Discrimination:
Tweets including mentions of exclusion, name-calling, or structural oppression were annotated as discriminatory (D1), if not then marked as (D0).
Example (annotation D1): "*When I had not announced my disability on a loudspeaker, I had some pretty awkward job interviews where they didn't know what to tell me. It made me understand why they add things like walking, lifting as functional requirements in central govt exams. They don't really want you.*"
For tweets marked as discriminatory, annotators further distinguished between personal accounts of discrimination and discrimination on a generic/societal level. The former was tagged as (DS) and the later as (DG).

Harassment:
Tweets related to disability-related harassment, including bullying, trolling, or abuse on a personal or societal level, were marked as harassment (H1) or (H0) if not.
Example (annotation: H1): "*If you want to know the social status of persons with disability in India, you should see conversations on reservations on social media. The use of words handicapped, viklang, not just in literary terms will reveal a lot to you.*"

Similar to Discrimination, we annotated personal accounts of harassment as (HS) and generic accounts of harassment as (HG).
After this, we also annotated Sexual harassment (Sexual) and Other (Other). These annotations were only applied to tweets already marked as harassment (H1).
Example (annotation: Sexual): ''*disabled face sexual abuse, domestic violence forced sterlisation, .its so diff for them to combat wethepeople loveknowsnodisability*"

Inclusion:
Inclusion-related tweets indicate positive experiences with accessibility, such as being able to make use of accommodations or witnessing thoughtful media representation. Such tweets were marked as (I1), or else (I0).
Example (annotation I1): "*Excited to find this watch with dial for persons with impairments*"

Advocacy:
Tweets related to disability-related advocacy, such as those calling for the rights or inclusion of people with disabilities on a personal or societal level, were marked as (A1), otherwise (A0).
Example (annotation A1): "*Accessibility modifications are required to enable persons with reduced mobility to gain access to education, employment, transportation*"
Similar to Discrimination, we annotated Self-Advocacy (AS) and Generic Advocacy (AG).

Identity:
If the tweet author referred to their own identity as a disabled person within the text of the tweet, it was annotated as identity (Id1), otherwise (Id0).
Example (annotation: Id1): "*That satisfying moment when,as a blind lawyer at a firm,you get to speak for work w/ a fellow blind lawyer who is your client.*"

Fact or Opinion:
If the tweet included factual and verifiable information, such as that about a government policy, cited statistics, court statement, or cited experiences from published articles, we marked it as factual (F). If the tweet included non-factual information such as opinions, personal experiences, or commentary on the state of disability, we marked it as opinion (O).
Example (annotation: F): "*freedom of a woman to decide whether to continue with a pregnancy cannot be taken away, the Kerala High Court has said while allowing a woman with multiple disabilities to abort*"
Example (annotation: O): "*Self- Care tip! I experienced panic attack and anxiety and I understood how important it is to take care of ourselves and reach out for help.*"

Stance:
If the author's stance on the issue described in the

tweet was positive, relatively positive, or hopeful, annotators were asked to mark it as positive (Pos). If the stance was negative, relatively negative, or critical, it was marked as negative (Neg).
Example (annotation: Pos): "*It is hoped that in months streets not less than km each in the South, East, North, West; Central Delhi will be identified; made accessible under supervision of an officer of a rank not lower than the Director be appointed by the Chief Secretary*"
Example (annotation: Neg): "*While I want to correct everyone who is saying 'specially' abled child while talking about the Ranchi airport incidence, I guess 'special' generates more empathy!*"

<u>**Theme**</u>:
Tweets related to Health and Hygiene, Education, and Employment were annotated as (HH), (Ed), and (Emp), respectively. Tweets that did not fall into these categories were annotated as Other (O).
Example (annotation: HH): "*What about inclusive accessible toilets for people with disabilities? Why not have unisex inclusive accessible toilet for both disabled &; trans people? Do frame EOP mandated u/s of too*"
Example (annotation: Ed): "*I am an aspiring deaf woman (1st in country) pursuing LLB in Faridabad. It is ironic how while learning to advocate for Deaf Rights, I've to struggle for my right to Interpreter provision! Pl support my quest for access to education!*"
Example (annotation: Emp): "*Working in Banking sector is getting difficult day by day, planning to quit as soon as possible. I know being visually impaired it will be difficult to get a new job especially when you have passed around years there but I will have to take risk. I feel suffocated now.*"

5 Annotation and Quality Control

We manually annotated the tweets to provide a solid benchmark and foster future research. The first two authors of the paper went through a pilot annotation exercise to verify the quality of their annotation schema and guidelines along with two other annotators. For the pilot study, we sampled 250 Tweets from our collection following the criteria: 1) the sample contains a considerable percentage of tweets containing disability related keywords and 2) some of the tweets are related to employment, education and health, and 3) the rest of the sample consists of random tweets not related to the above topics. The annotation is based only on the actual text of the tweet without considering additional modalities (e.g. images). This is similar to the information available to the predictive models at the time of training. After the first round of annotations, the inter-annotator agreement was calculated with a pairwise comparison between the annotators using Fleiss's Kappa (κ) for all the categories. Figure 2 lists the agreement values for each annotation category. Overall, high inter-rater reliability scores were achieved over all categories.

<u>**Adjudication**</u>: The last step of the pilot annotation was to reconcile disagreements among the annotators to produce the final canonical annotation. This step also allowed us to further refine the annotation guidelines. For example, whether a tweet is a fact or an opinion could sometimes be ambiguous and the annotators had to carefully consider and decide whether or not a user was stating opinions as facts. As a result, we refined the definition of "facts" to clearly include a condition that it belong to a set which is universally true. Take the following Tweet as an example: "*people who live in places which have free healthcare are privileged. just saying.*" This Tweet is a classic example of the user's opinion being stated as a fact. But since this statement is not universally true, we classified it as an opinion.

<u>**Main Annotation**</u>: Following the pilot, each annotator annotated mutually exclusive set of tweets. The annotators who designed the schema (average Cohen's Kappa across all the categories = 0.81) annotated 1,600 tweets between them, while the remaining 784 tweets were annotated by two other annotators. The average Fleiss' Kappa for all annotators over all the categories was 0.70, indicating high agreement. Table 2 shows the high-level statistics of the annotation of 2,384 tweets in the dataset.

6 Benchmarking Experiments

<u>**Text pre-processing:**</u> We pre-processed the tweets using TweetPreprocessor API[2] which helps in cleaning the tweet by parsing URLs, Hashtags, Mentions and Emojis.

<u>**Classification Models:**</u> We designed the annotation schema in a way that the majority of categor-

[2]https://pypi.org/project/tweet-preprocessor/

Classifiers	TF-IDF+LR			Bert-Base-Cased			RoBERTa-Base			BERTweet		
	Prec	Rec	F1	Prec	Rec	F1	Prec	Rec	F1	Prec	Rec	F1
BC (Relatedness)	0.71	0.83	0.76	0.86	0.83	0.84	0.91	0.95	0.93	0.88	0.95	0.91
MC.1 (Theme)	0.45	0.67	0.54	0.65	0.72	0.68	0.67	0.77	0.72	0.70	0.82	0.76
BC.1 (Discrimination)	0.45	0.56	0.50	0.64	0.71	0.67	0.69	0.80	0.74	0.78	0.85	0.81
BC.2 (Advocacy)	0.57	0.50	0.53	0.52	0.53	0.52	0.61	0.43	0.51	0.58	0.55	0.56
BC.3 (Identity)	0.62	0.65	0.63	0.70	0.81	0.75	0.66	0.95	0.82	0.69	1.00	0.81
BC.4 (Harassment)	0.53	0.56	0.54	0.55	0.60	0.58	0.59	0.61	0.60	0.63	0.66	0.64
BC.5 (Inclusion)	0.62	0.63	0.62	0.64	0.65	0.64	0.66	0.63	0.64	0.69	0.67	0.68
BC.6 (Fact/Opinion)	0.61	0.45	0.56	0.62	0.54	0.58	0.67	0.55	0.6	0.72	0.56	0.66
BC.6.1 (Stance)	0.85	0.88	0.86	0.84	0.77	0.80	0.95	0.83	0.88	0.96	0.85	0.90

Table 3: Classification Report using different bag-of-words and transformer models on the test set of the annotated dataset. Prec indicates Precision, Rec indicates Recall and F1 indicates F1-score averaged over all the class labels.

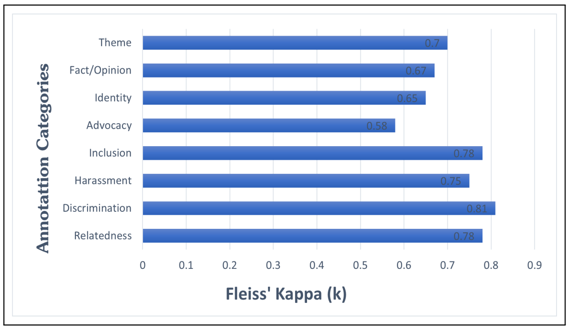

Figure 2: Inter-Annotator Agreement (κ) among the annotators on the categories.

ical themes, such as Discrimination or Not Discrimination, can be determined using binary classification. In contrast to multi-class hierarchical classifiers, such binary classifiers do not require a large amount of training data. We therefore took the approach of developing separate classifiers for tagging each category.

The top-most level **(Level-1)** used a binary classifier (BC) to determine whether the tweet is related to disability **(BC)**. If the output of **(BC)** was 'Yes', we then used six different binary classifiers in the second level of tagging **(Level-2)** to determine if the tweet was related to 1) *discrimination* **(BC.1)**, 2) *advocacy* **(BC.2)**, 3) *identity* **(BC.3)**, 4) *harassment* **(BC.4)**) *inclusion* **(BC.5)**, or if it was a 6) fact or opinion **(BC.6)**. Moreover, in Level-2, we also designed a multi-class multi-label classifier to examine the domain or theme the tweet pertains to, for example, employment, education, health or others **(MC.1)**. Based on the outputs obtained from Level-2 classification, we designed six binary classifiers **(Level-3)** to examine if the discrimination was self-experienced or generic **(BC.1.1)** if the output was 'Yes', similarly for advocacy **(BC.2.1)** if the output was 'Yes', harassment **(BC.4.1)** if the output was 'Yes', stance if the tweet was opinionated **(BC.6.1)** . We also designed another classifier to determine the nature of the harassment **(BC.4.2)**.

Training and Evaluation: Each of the classifiers were separately trained on class-balanced training data for each annotation category (such as binary classification to determine discrimination). We trained each model three times using different random seeds and reported the mean Precision, Recall and F1 (macro) on the test set. For all the annotation categories (binary and multi-class classification), we benchmarked the dataset using the following baselines:

TF-IDF+LR: We trained a Logistic Regression (LR) with the TF-IDF vectors of the input tweets using L2 regularization.

BERT, RoBERTa and BERTweet: We evaluated the vanilla transformer-based models (Vaswani et al., 2017), such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019) and BERTweet (Nguyen et al., 2020) from huggingface Transformers[3]. BERTweet is pretrained on English tweets using RoBERTa as the encoder and it achieves better performance on Twitter tasks (Nguyen et al., 2020). We fine-tuned the BERT, RoBERTa and BERTweet for binary (BC, BC.1, BC.2, BC.3, BC.4, BC.5, BC.6, BC.1.1, BC.2.1, BC.4.1, BC.4.2, BC.6.1) and multi-class (MC.1) predictions by adding a classification layer that took the [CLS] token as input. We used the base cased models and fine-tuned them for 10 epochs. The maximum sequence length was set to 50 in the

[3]https://huggingface.co/docs/transformers/index

Advocacy		Discrimination	
Unigrams	**Bigrams**	**Unigrams**	**Bigrams**
champion	accessible screen	dying	struggle invisible
deaf	delhi govt	unable	hide vulnerabilities
raised	high support	miserably	hide divyang
girlspl	home vaccination	unfortunately	people comorbidities
high	support needs	humiliating	visually impaired
Inclusion		**Identity**	
inclusive	accessible flight	us	instant intimidation
excited	disabled friendly	great	hearing aid
accepted	accessible india	blind	visually impaired
included	application accessible	deaf	blind woman
accessibile	education accessible	flag	deaf woman
Harassment		**Themes**	
violent	getting beaten	vaccination	educate disabled
abuse	disabled women	reservation	disabled friendly
deaf	disabled unfriendly	universities	home vaccination
flag	home vaccination	covid	educational institutions
marry	support needs	employment	education system

Table 4: Top 5 Unigrams and Bigrams Association in case of Discrimination, Advocacy, Harassment, Inclusion, Identity and Accessibility theme sorted by Pearson Correlation. All correlations are significant when considering $p < .01$ determined using two-tailed t-test.

training set and used a batch size of 32.

Experimental Results: Table 3 shows the predictive performance of all the models for the different categories (i.e., both binary and multi-class classification). Overall, BERTweet models with linguistic information achieved better overall performance. Transformer models performed substantially better in the majority class baseline and above Logistic Regression. BERTweet performed better than BERT and RoBERTa, which illustrates the advantage of pre-training on English tweets for this task. These results indicate that the transformer models achieve acceptable predictive performance on categories, such as *Relatedness, Theme, Discrimination, Inclusion, Identity, Stance*. However, it is evident that there is much room for improvement for classifiers on categories, such as *Advocacy, Inclusion, Harassment and Fact/Opinion* as they considerably under-perform compared to human judgement.

7 N-gram Analysis

To understand the most prominent and distinguishing patterns in each category, we used unigram and bigram tags associated with the annotated categories of the tweets in our data set. Each tweet was represented as a TF-IDF distribution over the unigrams and bigrams to reveal distinctive syntactic patterns of different categorical themes. For each feature, we computed the strength of correlation between its distribution across posts and the label of the post using Pearson Correlation (r) (Benesty et al., 2009) – a standard approach used by other researchers (Jin et al., 2022). Finally, we sorted these values and obtained the most important n-grams for each category.

Table 4 presents the top 5 unigrams and bigrams correlated with our six annotation categories. The top n-grams in the harassment and discrimination category can be classified into (a) negative verbs and adjectives (e.g. *violent, deaf, getting beaten, disabled unfriendly, humiliating*) that usually depict the kind of societal harassment disabled people in India experience in their everday lives; and (b) word spans related to the trend of reacting to harassing or discriminatory experiences (e.g. *hide vulnerabilities, hide divyangs*[4]).

On the other hand, the most important features in advocacy/inclusion categories can be classified into positive nouns and supportive or encouraging keywords (e.g. *champion, accepted, high support, accessible india*); and (b) some suggestions on improving access to vaccination, transportation (e.g. *application accessible, education accessible*).

In the identity category, we observe that most n-grams are related to people disclosing their status as a disabled person. Similarly, in the themes, there is a high degreee of association in education and employment related keywords. One interesting highly frequent n-gram is "reservation", and it appears that disabled people are vocal about affirmative action in education and employment.

8 Analysis on Disability and Gender

Since gender and professions play a crucial role in shaping up the ways in which people express themselves on social media, we conducted a preliminary quantitative and content analysis on our corpus to determine the gendered differences in patterns of self-expression of people with disabilities on Twitter. We illustrate three preliminary observations emerging from our analysis:

1. Disabled female users center personal experiences while tweeting about discrimination, advocacy and harassment more frequently than disabled male users.

Figure 4 shows that 15% of the tweets from the male handles were on discrimination, 12% were on advocacy of rights and 10% were *personal* accounts of harassment. In contrast, 22% of the

[4]"Divyang" is a Hindi-word meaning disabled.

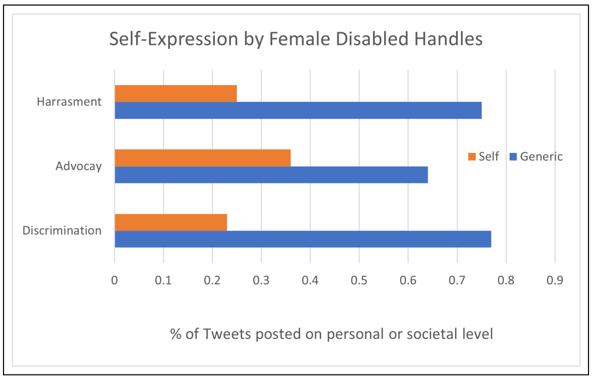

Figure 3: Pattern of Self-Expression by Female Disabled Handles on Indian Twitter.

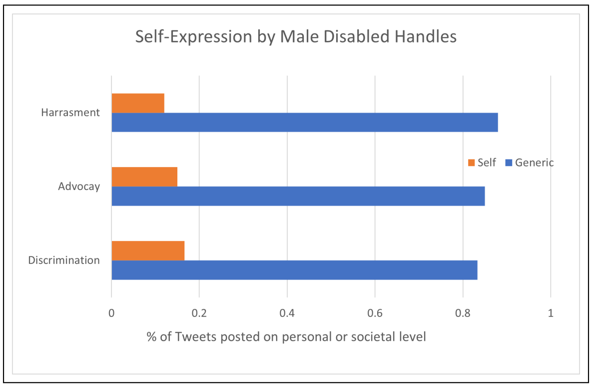

Figure 4: Pattern of Self-Expression by Male Disabled Handles on Indian Twitter.

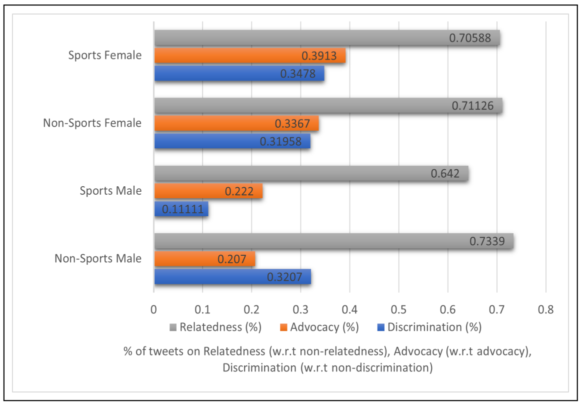

Figure 5: Pattern of Self-Expression by disabled women and men in Sports and Non-Sports on Twitter.

tweets from the female handles described discrimination related issues, 35% advocated for rights for the disabled and 26% described *personal accounts* of harassment. We perform a statistical significance test using Chi-Square (McHugh, 2013) to determine the gendered differences between the patterns of expressing advocacy, harassment and discrimination. The p-values obtained were 0.019, 0.599 and 0.011 for advocacy, harassment and discrimination patterns, respectively. Except harassment, the other values were statistically significant.

The content analysis revealed that male users are more likely to comment on broader, structural issues underlying discrimination, such as exclusionary government policies. Female users, on the other hand, publish a larger number of tweets centering personal experiences of exclusion and disability-related discrimination. Disability studies work often cites that disability – associated with being ‘dependent and helpless’ – is in conflict with masculinity, which is associated with being powerful and autonomous (Shuttleworth et al., 2012). Within India’s deeply patriarchal society in which ableist norms stigmatize male expression of need, this ‘masculinity dilemma’ may disincentivize male users from sharing personal experiences on public profiles. Given a disproportionate burden of discrimination, disabled women may have simply have a larger bank of discrimination-related personal experiences to draw upon. Further, within economies of visibility, highly visible women are more likely to perform the labour of authenticity (Duffy, 2015; Toffoletti and Thorpe, 2018; Banet-Weiser, 2021). Since there is a link between personal vulnerability and online harassment (Duffy and Hund, 2019), this opens up an avenue for further research on experiences of disabled women in India with online harassment.

2. Female paralympians publish positive tweets on inclusion more frequently than disabled women in other professions as well as men in all professions.

Figure 5 shows the quantitative distribution pattern which indicates that disabled sportswomen play a much larger role in tweeting about inclusion (39% of the tweets) compared to disabled women in other professions (33% of the tweets). The difference in advocacy patterns between males with disabilities in sports (22%) and those in other professions (21%) is marginal.

Our content analysis shows that while Paralympians tweet about inclusion, they often use positive tonality, praising the government for new policies, schemes, and initiatives. They also receive significant media engagement from political influencers and government bodies (French and Le Clair, 2018; Mitchell et al., 2021;

Pate et al., 2014; Toffoletti, 2018). Previous work shows that Indian sportspeople tend to use Twitter to support the government (Mishra et al., 2021) – a phenomenon rooted in the State's attempts to garner political support from influential figures. This celebration of disabled people in sports is part of the creation of a national identity centered around empowerment and unity. However, the disabled body is positioned as a form of 'apolitical diversity' – a condition produced by the conflation of nationalism and neoliberalism (Friedner, 2017). In such cases, Paralympians may come to be constructed as inspirational *'feel-good'* figures who are disincentivized from appearing to be critical online. This finding also points to the fact that online performances of positivity themselves may be gendered among influential disabled users. We note that a marginal percentage of Paralympians in our set acted against this norm, tweeting about non-reception of promised rewards, such as jobs and monetary payouts for achievements in Paralympic sports. This is a valuable insight showing that disability-related discrimination in India is the norm for even the most influential figures.

3. Disabled women are less vocal about facing harassment than disabled men.

From the distribution of tweets generated by disabled men, we found that 18% of male users raised their voices about harassment either on a broad *societal or personal level*, whereas the percentage of disabled women doing the same was only 6%.

There is overwhelming evidence that disabled Indian women face disproportionately more harassment in contrast to disabled Indian men. That online self-expression is not reflective of this points not only to the perceived stigma of mentioning harassment on Twitter, but that for women, many such discussions may occur in private online communities rather than the public sphere of participatory social media. Previous work has also shown that Indian women often limit self-expressions on topics intersecting with patriarchy (Karmakar, 2021).

9 Conclusion

This paper introduced a novel human-annotated corpus "#**DisabledOnIndianTwitter**" comprising of tweets posted by disabled people in India from a diverse set of professions. We manually tagged the corpus to categorize different patterns of self-expression based on a hierarchical annotation taxonomy. Using our corpus, we next conducted quantitative and content analysis to identify gendered differences in expressions of disabled people on Indian Twitter. We believe that the annotation schema as well as the dataset can be valuable in understanding social media use by disabled people. We aim to make our dataset publicly available to foster research at the nexus of NLP and Accessibility.

10 Ethics Statement

The use of Twitter data for research purposes is subject to the Developer Policy and Agreement. In accordance, aggregate analysis of Twitter content, including that related to sensitive topics such as health, that does not store any personal data, is permitted (Twitter). We followed these guidelines and stripped our data of user IDs, usernames, and other identifiers in order to protect the anonymity of users. Our set only includes tweets published in the *public* domain, by users who disclosed their disabled identity in their Twitter bio, profile picture, username, display name, or within the content of their tweets. In this way, we attempt to avoid making assumptions about the status of users' disabilities.

References

Marzieh Babaeianjelodar, Stephen Lorenz, Josh Gordon, Jeanna Matthews, and Evan Freitag. 2020. *Quantifying Gender Bias in Different Corpora*, page 752–759. Association for Computing Machinery, New York, NY, USA.

Sarah Banet-Weiser. 2021. Gender, social media, and the labor of authenticity. *American Quarterly*, 73(1):141–144.

Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2018. Fairness and machine learning limitations and opportunities.

Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. Pearson correlation coefficient. In *Noise reduction in speech processing*, pages 1–4. Springer.

Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Tauman Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *NIPS*.

Alexis Buettgen, Rachel Gorman, Marcia Rioux, Kuhu Das, and Soumya Vinayan. 2015. Employment, poverty, disability and gender: A rights approach for women with disabilities in india, nepal and bangladesh.

Census. 2011. *Measurement of Disability through Census National Experiences*.

Alison Croft. 2013. Promoting access to education for disabled children in low-income countries: Do we need to know how many disabled children there are? *International Journal of Educational Development*, 33(3):233–243. Educating the poorest.

Ranjita Dawn. 2014. "our lives, our identity": women with disabilities in india. *Disability and rehabilitation*, 36(21):1768–1773.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Bishnu Maya Dhungana. 2006. The lives of disabled women in nepal: vulnerability without support. *Disability & Society*, 21(2):133–146.

Emily Dinan, Angela Fan, Adina Williams, Jack Urbanek, Douwe Kiela, and Jason Weston. 2019. Queens are powerful too: Mitigating gender bias in dialogue generation. *CoRR*, abs/1911.03842.

Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '18, page 67–73, New York, NY, USA. Association for Computing Machinery.

Brooke Erin Duffy. 2015. Gendering the labor of social media production. *Feminist Media Studies*, 15(4):710–714.

Brooke Erin Duffy and Emily Hund. 2019. Gendered visibility on social media: Navigating instagram's authenticity bind. *International Journal of Communication*, 13:20.

Lata Dyaram and Suresh V. 2020. Workplace disability inclusion in india: review and directions.

Katie Ellis and Gerard Goggin. 2013. Disability and social media. In *The social media handbook*, pages 134–151. Routledge.

Liam French and Jill M Le Clair. 2018. Game changer? social media, representations of disability and the paralympic games. In *The Palgrave handbook of paralympic studies*, pages 99–121. Springer.

Michele Friedner. 2017. How the disabled body unites the national body: disability as 'feel good'diversity in urban india. *Contemporary South Asia*, 25(4):347–363.

Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H. Chi, and Alex Beutel. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, page 219–226, New York, NY, USA. Association for Computing Machinery.

Nora E Groce, Gayatri Kembhavi, Shelia Wirz, Raymond Lang, Jean-François Trani, and Maria Kett. 2011. Poverty and disability – a critical review of the literature in low and middle-income countries. *WGSRN: Other Gender & Development (Sub-Topic)*.

Saad Hassan, Matt Huenerfauth, and Cecilia Ovesdotter Alm. 2021. Unpacking the interdependent systems of discrimination: Ableist bias in NLP systems through an intersectional lens. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3116–3123, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Bronwyn Hemsley, Stephen Dann, Stuart Palmer, Meredith Allan, and Susan Balandin. 2015. "we definitely need an audience": experiences of twitter, twitter networks and tweet content in adults with severe communication disabilities who use augmentative and alternative communication (aac). *Disability and rehabilitation*, 37(17):1531–1542.

Ben Hutchinson, Vinodkumar Prabhakaran, Emily Denton, Kellie Webster, Yu Zhong, and Stephen Denuyl. 2020. Social biases in NLP models as barriers for persons with disabilities. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5501, Online. Association for Computational Linguistics.

Jens Ineland, Stefan Gelfgren, and Coppélie Cocq. 2019. Negotiating authority: Disability, interactions and power relations on twitter. *Scandinavian Journal of Disability Research*, 21(1):238–249.

Syed Salma Jameel. 2011. Disability in the context of higher education: Issues and concerns in india. *Electronic Journal for Inclusive Education*, 2:3.

Sophie Jentzsch, Patrick Schramowski, Constantin Rothkopf, and Kristian Kersting. 2019. Semantics derived automatically from language corpora contain human-like moral choices. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '19, page 37–44, New York, NY, USA. Association for Computing Machinery.

Mali Jin, Daniel Preotiuc-Pietro, A. Seza Doğruöz, and Nikolaos Aletras. 2022. Automatic identification and classification of bragging in social media. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3945–3959, Dublin, Ireland. Association for Computational Linguistics.

Aasha Kapur Mehta and Amita Shah. 2003. Chronic poverty in india: Incidence, causes and policies. *World Development*, 31(3):491–511. Chronic Poverty and Development Policy.

Priyanka Karmakar. 2021. Missing daughters: Social perceptions and treatment of the girl child in india.

Arun Kumar, Deepa Sonpal, and Vanmala Sunder Hiranandani. 2012. Trapped between ableism and neoliberalism: Critical reflections on disability and employment in india. *Disability Studies Quarterly*, 32.

S. G. Leveille, H. E. Resnick, and J. Balfour. 2000. Gender differences in disability: Evidence and underlying reasons. *Aging Clinical and Experimental Research*, 12(2):106–112.

Shahar Levy, Koren Lazar, and Gabriel Stanovsky. 2021. Collecting a large-scale gender bias dataset for coreference resolution and machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2470–2480, Punta Cana, Dominican Republic. Association for Computational Linguistics.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Kelly Mack, Emma McDonnell, Dhruv Jain, Lucy Lu Wang, Jon E. Froehlich, and Leah Findlater. 2021. What do we mean by "accessibility research"? a literature survey of accessibility papers in chi and assets from 1994 to 2019. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA. Association for Computing Machinery.

Islay Mactaggart, Hannah Kuper, G. V. S. Murthy, Jayanthi Sagar, Joseph Oye, and Sarah Polack. 2016. Assessing health and rehabilitation needs of people with disabilities in cameroon and india. *Disability and Rehabilitation*, 38(18):1757–1764. PMID: 26680511.

Benjamin W Mann. 2018. Rhetoric of online disability activism:# cripthevote and civic participation. *Communication Culture & Critique*, 11(4):604–621.

Mary L McHugh. 2013. The chi-square test of independence. *Biochemia medica*, 23(2):143–149.

Nilika Mehrotra. 2021. Emergent Disability voices on Social Media during COVID -19 times.

George A. Miller. 1995. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41.

Dibyendu Mishra, Ronojoy Sen, and Joyojeet Pal. 2021. Sporting the government: Twitter as a window into sportspersons' engagement with causes in india and usa. *arXiv preprint arXiv:2109.07409*.

Fallon R Mitchell, Paula M van Wyk, and Sara Santarossa. 2021. Curating a culture: The portrayal of disability stereotypes by paralympians on instagram. *International Journal of Sport Communication*, 14(3):334–355.

Daniel Mont, Natasha Layton, Louise Puli, Shivani Gupta, Abner Manlapaz, Kylie Shae, Emma Tebbutt, Irene Calvo, Mahpekay Sidiqy, Kudakwashe Dube, and Ulamila Kacilala. 2021. Assistive Technology during the COVID-19 Global Pandemic: The Roles of Government and Civil Society in Fulfilling the Social Contract. *International Journal of Environmental Research and Public Health*, 18(22):12031.

Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. Bertweet: A pre-trained language model for english tweets. *arXiv preprint arXiv:2005.10200*.

Itto Outini. 2020. The Impact of Social Media on the Social Lives of People with Visual Impairment (Facebook Groups as a Case Study). *Graduate Theses and Dissertations*.

Joyojeet Pal, Ana Maria Huaita Alfaro, Tawfiq W. Ammari, Sidharth Chhabra, and Meera Lakshmanan. 2017. Representation, Access and Contestation: Facebook and Vision Impairment in Jordan, India, and Peru. Accepted: 2017-02-15T11:54:26Z Publisher: The Critical Institute.

Joshua R Pate, Robin Hardin, and Brody Ruihley. 2014. Speak for yourself: analysing how us athletes used self-presentation on twitter during the 2012 london paralympic games. *International Journal of Sport Management and Marketing*, 15(3-4):141–162.

Mónica Pinilla-Roncancio, Islay Mactaggart, Hannah Kuper, Carlos Dionicio, Jonathan Naber, G.V.S. Murthy, and Sarah Polack. 2020. Multidimensional poverty and disability: A case control study in india, cameroon, and guatemala. *SSM - Population Health*, 11:100591.

Shanya Sharma, Manan Dey, and Koustuv Sinha. 2021. Evaluating gender bias in natural language inference. *CoRR*, abs/2105.05541.

Russell Shuttleworth, Nikki Wedgwood, and Nathan J Wilson. 2012. The dilemma of disabled masculinity. *Men and masculinities*, 15(2):174–194.

Maya Thomas and Maliakal Joseph Thomas. 2002. Status of women with disabilities in south asia. *Selected readings in community-based rehabilitation, series*, 2.

Kim Toffoletti. 2018. Sport, postfeminism and women with disabilities: Female paralympians on social media. In *New Sporting Femininities*, pages 253–275. Springer.

Kim Toffoletti and Holly Thorpe. 2018. Female athletes' self-representation on social media: A feminist analysis of neoliberal marketing strategies in "economies of visibility". *Feminism & Psychology*, 28(1):11–31.

Twitter. More on restricted use cases – twitter developers.

Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. SemEval-2018 task 3: Irony detection in English tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 39–50, New Orleans, Louisiana. Association for Computational Linguistics.

Aditya Vashistha, Edward Cutrell, Nicola Dell, and Richard Anderson. 2015. Social Media Platforms for Low-Income Blind People in India. In *Proceedings* *of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*, ASSETS '15, pages 259–272, New York, NY, USA. Association for Computing Machinery.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Zeerak Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on Twitter. In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.

Jieyu Zhao, Subhabrata Mukherjee, Saghar Hosseini, Kai-Wei Chang, and Ahmed Hassan Awadallah. 2020. Gender bias in multilingual embeddings and cross-lingual transfer. *CoRR*, abs/2005.00699.

11 Appendix

12 Disability Related Keywords

"deaf", "mute", "blind", "one legged","disabled", "disability", "handicap", "crippled", "low vision", "visually impaired", "Hearing impairment", "Locomotor disability", "Attention Deficit Hyperactivity Disorder", "ADHD", "Muscular Dystrophy", "Hard of Hearing", "Parkinson's Disease", "dwarf", "short stature", "accessibility", "braille", "sign language", "autism", "dyslexia", "dysgraphia", "dyscalculia", "dyspraxia", "aphasia", "dysphagia", "multiple sclerosis", "cerebral palsy", "genetic disorders", "arthritis", "heart failure", "insanity", "mental illnesses", "depression", "bipolar disorder", "paralysis", "wheelchair", "hearing aid", "epilepsy", "chronically ill", "down's syndrome", "retard", "Asperger Syndrome", "Alzheimer's" .

Gender	#Followers	#Tweets	Mention of Disability/Profession in Bio
Female	34.4k	6308	Para badminton player
Female	81k	977	Bomb Blast Survivor
Female	19k	1699	Amputee climb Mt Everest
Female	61.6k	6445	Paralympian
Female	1055	14400	Crip, queer artist, consultant
Female	270	1780	Author ǀ Blogger
Female	7363	576	Deaf chess champion
Female	267	934	Gay, resentful
Female	309	996	Lifestyle blogger
Female	23.1k	28700	Autistic actor
Female	128	765	Law Student
Female	70	972	Mrs India 2021
Female	1055	14400	Crip, queer artist, consultant
Female	270	1780	Author ǀ Blogger
Female	890	56500	Lawyer, Comedian
Female	407	853	Traveller
Female	71	42	Disability Inclusion Facilitator
Female	11.3k	1400	Managing Director,@Jindal$_SAW$
Female	521	3897	activist, comedian, writer
Female	388	61	Paralympian
Female	49k	235	Paralympian
Female	62	94	Indian Para Athlete
Female	1245	762	researcher, artist, and author
Female	629	390	Chief Content Officer
Female	403	136	Deaf Woman pursuing Law
Female	193	67	International Tennis Player
Female	865	1828	Aspiring Biologist
Male	1062	4253	Writer. Poet. Disabled.
Male	7437	25900	Disability Rights Defender
Male	5992	6445	Indian Para Swimmer
Male	151	724	an atheist, fan of test cricket
Male	2604	3364	Lawyer, Rhodes Scholar
Male	387	545	Para Archer
Male	28.5k	600	Javelin Thrower Paralympic
Male	96	507	Professor, Research Scholar
Male	80	51	deaf, Indian sign language
Male	51	71	Deaf Postal Assistant
Male	110	63	Deaf
Male	150	254	lawyer
Male	524	2357	L-Vision,(Blind) student
Male	750	10400	Deaf journalist
Male	197	530	visually impaired athlete

Table 5: Details of the Disabled Twitter Handles of India considered for our study.