



MD3: The Multi-Dialect Dataset of Dialogues

Jacob Eisenstein¹, Vinodkumar Prabhakaran¹, Clara Rivera¹
Dorotya Demszky², Devyani Sharma³

¹Google Research, ²Stanford University, ³Queen Mary University of London

{jeisenstein, vinodkpg, rivera}@google.com
ddemsky@stanford.edu, d.sharma@qmul.ac.uk

Abstract

We introduce a new dataset of conversational speech representing English from India, Nigeria, and the United States. The Multi-Dialect Dataset of Dialogues (MD3) strikes a new balance between open-ended conversational speech and task-oriented dialogue by prompting participants to perform a series of short information-sharing tasks. This facilitates quantitative cross-dialectal comparison, while avoiding the imposition of a restrictive task structure that might inhibit the expression of dialect features. Preliminary analysis of the dataset reveals significant differences in syntax and in the use of discourse markers. The dataset, which will be made publicly available with the publication of this paper, includes more than 20 hours of audio and more than 200,000 orthographically-transcribed tokens.¹

Index Terms: dialect, world Englishes, dialogue


1. Introduction

A key research challenge for spoken language processing is to build systems that meet users where they are: in their own languages and dialects. While there has been significant progress towards *multilingual* speech and text processing (e.g., [1]), the development of *multidialectal* systems, datasets, and evaluations lags behind. Because billions of people speak dialects of global languages such as English, Arabic, and Spanish, multidialectal speech processing could dramatically increase access to language technology.

In this paper, we present MD3, the **Multi-Dialect Dataset of Dialogues**. The current release of MD3 includes a total of 20 hours of audio from three varieties of global English: India, Nigeria, and the United States. Unlike most previous datasets of dialectal speech, which focused largely on scripted speech (e.g., [2]) or open-ended conversations (e.g., [3]), MD3 is centered on information-sharing tasks with clearly-defined speaker intents. This makes it possible to study the dialect robustness of spoken-language processing systems not only in phonology but also in downstream language-processing tasks that are closely related to applications such as information retrieval and question answering. But unlike task-based dialogue scenarios, MD3 has no limitation on the vocabulary or grammatical structure, facilitating a more uninhibited style of interaction in which dialect features are more likely to appear [4].

The MD3 conversations are organized around guessing games, in which one speaker (the “describer”) must commu-

¹The MD3 dataset is available at <https://www.kaggle.com/datasets/jacobeis99/md3en>, under the CC BY-SA 4.0 license. Jon Clark played a key role in the early stages of this research. Many thanks to Sunny Mak and Ravi Rajakumar for the design and implementation of the elicitation software. For more details and acknowledgments, please see <https://arxiv.org/abs/2305.11355>.



D: Here is a dog image. It is in grey colour. Ah uh
→ um we can see one chain in his neck.
D: Background, we can see grass.
G: It is looking left side?
D: Ah it's looking left side but his face towards
→ camera only.
G: Okay, okay. Done.

Figure 1: An example dialogue from two en-IN speakers. In the transcript, ‘D’ indicates the describer, who sees only the single image shown in the upper left; ‘G’ indicates the guesser, who sees all twelve images.

nicate a piece of information to the other (the “guesser”). There are two types of games: an image-guessing game (Figure 1), in which the describer must describe an image well enough for the guesser to select it from a set of twelve similar images, and a word-guessing game (Figure 2), in which the describer must communicate a word or phrase while avoiding a list of related words. This methodology elicits speech that is comparable in register and topic, without telling the participants what to say. The current release includes 3689 such games, orthographically transcribed into approximately 200,000 words (see Table 1). We also release metadata about the guessing games that prompted each dialogue. We hope that this dataset will serve as a benchmark for dialect-robust spoken language processing and as a resource for the study of global English.

2. Related work

Early work on “accent classification” focused on datasets of short scripted speech [2]. Subsequent shared tasks on language identification included the classification of en-US and en-IN in spontaneous conversational telephone speech [5], using the CALLFRIEND corpora [3]. Multi-dialect datasets of conversational speech have been gathered in other languages, including Arabic [6], Austrian [7], Swiss German [8], and Spanish [9]. Beyond dialect classification, researchers have explored the detection of specific dialect features [10, 11] and the quantitative density of dialect features [12, 13, 14]. In the speech domain, such work has focused primarily on unconstrained narrative-

based corpora such as the International Corpus of English [15] and the Corpus of Regional African American Language [16].

While unconstrained speech data has been a powerful resource for the study of dialect variation and the development of dialect-robust speech *recognition*, it is less ideal for the development of dialect-robust *speech processing systems*. This genre of speech often touches on local entities, such as place names, which are strong signals to the locale of the conversation while carrying little interesting dialectological information [10]. Furthermore, unconstrained conversation is not directly applicable to the typical tasks of interest for speech processing systems, such as semantic parsing and question answering. The MD3 corpus addresses the first issue by restricting conversation to a predetermined set of topics governed by the information-sharing games that were used as conversational prompts. Regarding applicability to speech processing tasks, we have tried to strike a middle ground between unconstrained conversation and task-oriented dialogue systems (e.g., [17]), which are likely to inhibit dialect features by imposing a rigid task model. The information-sharing tasks in MD3 are not derived from any specific application, but are related to challenging queries in search and image retrieval. Most importantly, accuracy can be measured at the level of individual prompts, making it possible for users of the dataset to extend fairness analyses from prior work on speech recognition (e.g., [13, 18]) to downstream components of the speech processing pipeline.

3. Elicitation

The elicitation was performed in parallel in three locales: the United States (en-US), India (en-IN), and Nigeria (en-NG). In each locale, we constructed a pool of speakers from which we selected random pairs of individuals for a set of **matches**. No individual speaker participated in more than six matches and no pair participated in more than a single match. Each match was divided into five **rounds**, in which one speaker was given the role of **describer**, and the other speaker was given the role of **guesser** (see Figure 1). Within each round, the guesser and describer received a series of role-specific **prompts**, which required the describer to convey information to the guesser. Each round was five minutes long, and the participants received randomly-selected prompts until the time expired. The same pool of prompts was used in each locale.

The elicitation procedure was subject to an internal review, ensuring that we obtained informed consent from the participants and protected their privacy.

3.1. Speakers

Speakers were recruited by two third-party vendors. In each locale, the vendor recruited an equal number of female and male participants. All participants were at least 18 years old. Geographically, we targeted three broad and diverse geographical regions. To ensure a level of linguistic coherence within each region, we imposed additional demographic criteria:

- en-IN: native speaker of Telugu; high proficiency in English; recruited in Hyderabad.
- en-NG: native speaker of Yoruba; raised and educated in English.
- en-US: native speakers of English; born in the Western United States (U.S. Census region 4, district 9).

Demographic criteria were self-reported. There is considerable linguistic heterogeneity within each group: for example, the en-

Target Science Fiction

Taboo words Future, Imaginary, Advancements, Time, Travel

D: Emm do you know when you are in secondary school,
 ↪ you you are when you get to like ss1 there are
 ↪ three classes that you can be. It's either you
 ↪ are in, do you know those three classes?
 G: Commercial.
 D: Sorry?
 G: Commercial, Science, Art.
 D: Ok. You know the middle one you mentioned,
 G: which one?
 D: the one the middle one you mentioned
 G: science

Figure 2: An example dialogue from two en-NG speakers, with 'D' indicating the describer and 'G' indicating the guesser. The describer uses a Nigerian cultural reference (ss1) to introduce the term science, and they later solve the full prompt.

US group includes speakers of African-American English.

3.2. Prompts

Participants played two types of guessing games: an image-guessing game and a word-guessing game. In the image-guessing game, the guesser is shown twelve similar images (see Figure 1), one of which is also shown to the describer. The participants must discuss what they see until the guesser can identify the describer's image; if the guesser clicks on the correct image, the prompt is marked as a **win**, otherwise it is marked as a **loss**. The word-guessing game is similar to the popular game "Taboo": the goal is for the guesser to identify a given term known to the describer, who may not use that term nor any of a set of five "taboo terms." An example is shown in Figure 2. If the guesser states the target term, the prompt is marked as a win; if the describer accidentally uses the target term or one of the forbidden terms, it is marked as a loss. In the word game, we rely on self-reports for these outcomes. Participants were not given any incentive to successfully solve the prompts and had the option to **skip** any prompt. Nonetheless, as shown in Table 2, the participants solved most prompts successfully — bearing in mind that in the word-guessing game, results are based only on self-reports.

3.2.1. Image prompts

The image prompts are drawn from three public datasets: FoodX-251 [19], CalTech-UCSD Birds [20], and Stanford Dogs [21]. These datasets were chosen because they offer fine-grained image classes which are not easy to distinguish with a single word or phrase. From the Stanford Dogs dataset we showed the guesser twelve images of the same breed of dog; from the CalTech-UCSD Birds dataset we showed images of the same species of bird; and from FoodX-251 we showed images of the same type of food (e.g., falafel). As shown in Figure 1, this restriction forces the participants to describe several visual details of the image, including the spatial arrangement, background, and colors.

3.2.2. Word prompts

The prompts for the word guessing game require the target words and the corresponding forbidden words. Target words were drawn from a number of sources: (1) two publicly-

	speakers	dialogues	rounds	prompts	utterances	hours	tokens	WER
en-IN	27	46	134	1103	11856	6.37	62318	33.2
en-NG	39	44	124	957	11482	6.48	67237	38.6
en-US	38	37	152	1629	13235	8.94	86314	22.3

Table 1: Statistics of the dataset by locale. For details on the distinction between dialogues, rounds, and prompts, see Section 3. Tokens are counted by simple whitespace delimiting. For details on the word error rate calculation, see Section 4.1.

locale	game type	win	loss	skip
en-IN	image	90.6	9.4	0.0
	word	77.5	9.6	12.9
en-NG	image	86.1	13.3	0.5
	word	67.8	8.5	23.7
en-US	image	91.3	8.7	0.0
	word	83.7	4.1	12.0

Table 2: Results for prompts by locale and type. In the image game, a loss occurs when the guesser clicks on the wrong image; in the word game, a loss is when the describer accidentally uses the target term or one of the forbidden words.

available Taboo datasets: Tabouid [22] and “NS”,² (2) a curated list of popular personalities, actors, singers, athletes, buildings, and statues based on Wikipedia’s popular pages listing,³ (3) a list of the most frequent English words from Education First,⁴ (4) a manually-curated word list originally designed for the acquisition of text-to-speech data. From all the prompts, we removed entities that were unlikely to be widely known in all locales, names of living politicians, as well as offensive and sexual terms. For NS prompts, we used the forbidden words that were provided in the dataset. For prompts from all other sources, we elicited forbidden words through crowdsourcing. For each target word, three raters were asked to provide five most common words that they would use to describe it. We aggregated the responses by selecting the words that were most commonly chosen across raters, breaking ties manually.

3.3. Recording

At the time the dataset was recorded, COVID-19 restrictions made it impossible for participants to meet in person. Instead, they joined virtual meetings using Google Meet (they were asked to turn off their cameras), and simultaneously logged on to a crowdwork interface that presented the prompts.⁵ Most participants worked from their homes. Participants recorded their conversations using a proprietary web interface that securely stored and transmitted the audio. No special audio equipment was provided; most participants used their laptop microphones and speakers, but some used headsets. Many of the recordings contain background noise. We excluded audio files in which it was not possible to transcribe both participants. Audio is stored in 16-bit linear PCM encoding at 48 kHz in wav format.

²<https://github.com/nehasinha/Taboo/blob/master/assets/cards.csv>

³https://en.wikipedia.org/wiki/Wikipedia:Popular_pages

⁴<https://www.ef.com/wwen/english-resources/english-vocabulary/top-3000-words/>

⁵Before their first match, all participants attended a group training, where they were provided written guidelines describing the guessing games rules with an overview of the technologies to be used during the recording process, and engaged in a supervised training round.

3.4. Transcription

We transcribed a subset of audio recordings in which both participants were clearly audible, relying on crosstalk for one of the two participants. Orthographic transcription was performed by speech transcription professionals using audio files that were segmented by prompt. In nearly every case, transcription was performed by an individual from the same geographic locale as the speakers, e.g., en-NG speakers were transcribed by workers in Nigeria. Each transcription was reviewed by a second worker for accuracy.

3.5. Locale differences

While the form of the elicitation was identical across locales, there were significant differences in practice. Many of these differences were due to the fact that the speakers were communicating remotely from their own homes, using their own equipment. In Nigeria, power outages caused a number of matches to be cancelled, and internet access was slow and unreliable. In the U.S., there was less background noise in the recordings, and speakers may have had access to higher-quality microphones. For these reasons, although we recorded an identical number of dialogues in each locale, the final collection includes 30% more audio time in en-US than the other two locales (see Table 1).

A second area of difference relates to the prompts. Despite our efforts to identify prompts that would be solvable in all locales — as well as in Spanish and Portuguese-speaking locales that will be included in a future dataset release — some prompts were clearly more difficult for the non-US speakers. The issue was more significant in the word prompts, which were skipped nearly twice as often in the en-NG subcorpus as in en-IN and en-US (see Table 2). Differences were smaller in the image game: although the transcripts indicate that participants were sometimes unfamiliar with the types of food shown, this was rarely necessary to solve the game because in each prompt all of the candidate images were the same type of food.

The US-based speakers completed prompts at a higher rate than the other two locales (3.0 prompts / minute in the U.S., 2.9 in India, and 2.5 in Nigeria). This could be attributed to both of the above factors (familiarity and technological resources), as well as other possible causes such as English proficiency and cultural differences.

4. Dataset

Basic statistics of the dataset are shown in Table 1. Roughly the same number of dialogues were recorded in all three locales, but as noted above, a greater proportion of the en-US dialogues were of sufficiently high audio quality to transcribe.

4.1. Speech recognition

As a first test of the difficulty of the dataset for speech recognition, we applied the Whisper speech recognition system, using the SMALL.EN checkpoint [23]. Word error rates (WER) are

Feature	Meaning/form	en-IN	en-NG	en-US
<i>Having as stative progressive</i>	standard	1	1	2
	extended	30	11	0
Focus <i>only</i>	exclusive	193	31	42
	focus	23	0	0
Confirmation markers	<i>okay</i> [all]	1362	543	1734
	<i>got it</i> [all]	435	209	367
	<i>okay got it</i>	75	2	35
	<i>gotcha</i>	0	0	13
	<i>done</i>	219	6	1

Table 3: *Dialect feature counts. For more details, see Section 4.2.*

shown in the rightmost column of Table 1, ranging from 22.3 in en-US to 38.6 in en-NG. This indicates that the dataset is relatively difficult: in the Whisper paper, only two of fourteen datasets yield a higher error rate than the MD3 en-US subcorpus (CHiME6 and AMI-SDM1), and none has a higher error than the MD3 en-NG subcorpus. Potential causes for these differences include dialect sensitivity of the Whisper model as well as the different levels of audio quality in each locale.

4.2. Dialect features

Goal-oriented interactive tasks help to divert attention from the recording situation and encourage naturalistic speech. As a result, the MD3 subcorpora exhibit many natural dialect features that can support a wide range of fine-grained analysis of language variation. For example, the speech samples in en-NG and en-IN instantiate naturalistic usage of classic phonological characteristics of Nigerian English [24], particularly those associated with Yoruba speakers, and Indian English [25], particularly those associated with Telugu speakers. Similarly, there is extensive variation in lexical, morphosyntactic, and dialogic features in the three subcorpora. We illustrate this systematic diversity with a snapshot of three syntactic and dialogic features in MD3, each with a distinct predicted distribution across the three dialects. Table 3 shows counts for these features.

The first feature is the use of progressive *-ing* with extended stative meanings, e.g. *it’s actually having its tongue slightly out* [en-NG]. Extended use of progressive *-ing* has been attested robustly for Indian English speakers [26] and also for Yoruba English speakers at a lower rate [27], due to shared grammatical properties in the local languages (L1s) of the two regions. By contrast, it is not a feature of American English. This is precisely instantiated in the corpus: 30 of 31 uses of *having* in en-IN are extended stative uses. These are also attested in en-NG, but at a lower rate (11 instances). As expected, there are no instances in en-US.

The second feature is the use of *only* with non-contrastive, presentational focus meaning rather than exclusive meaning, e.g. *Is it a chocolate cake? Yes yes chocolate cake only* [en-IN]. Unlike extended progressives, the L1 source of focus *only* arises in Indian languages [28] but not in Yoruba, and so is only predicted for en-IN, not en-NG and en-US, which should pattern together. We find that 23 of 216 instances of *only* in en-IN are associated with the novel, non-contrastive presentational focus meaning. As predicted, neither en-NG (31 instances of *only*) nor en-US (42 instances of *only*) have any such usage. Here en-IN stands apart not only in terms of two distinct meanings, but also in the overall frequency of use of *only*, also due to the prevalence of pragmatic markers in Indian languages.

Finally, discourse markers that affirm shared knowledge or agreement [29] are very prevalent in the dataset given the nature of the task, e.g. *gotcha* [en-US]. These confirmation markers show fine-grained patterns of overlap and difference across varieties. All three subcorpora show high levels of use of *okay* and *got it*, but each also shows distinctive behaviors: en-IN includes 219 uses of *done* as an affirmation marker, with only 1 and 6 uses respectively in en-US and en-NG. En-US has 13 instances of *gotcha*, a form absent in the two other dialects. And notably, en-IN shows a much higher overall use of confirmation markers than the other dialects, while en-NG uses half the overall amount used in en-US.

These examples show the MD3 corpus captures meaningful dialect differences and can extend our understanding of these dialects through discourse-level dialect features. Though not examined here, each region also exhibits inter-speaker variation. The orderly variation present in the closely parallel speech samples of the MD3 corpus represents a unique new resource for both dialect-robust spoken language processing and for the analysis of global English varieties.

4.3. Limitations

While the dataset demonstrates meaningful dialectal variation, researchers should be cautious when drawing generalizations about the speech patterns of the represented dialects or their speakers. First, the dataset includes speakers with one of three first languages: Telugu (India), Yoruba (Nigeria) and English (US). This is only a small subset of first languages spoken in these countries. Many of the dialect features are specific to the first language of the speakers, and hence the linguistic patterns may not generalize to speakers with other first languages. Second, the recruitment process relied on third-party vendors, which may have introduced selection biases. Third, in the word guessing game, it was difficult to select prompts that worked equally well across locales: given the dominance of western culture in Wikipedia, it likely that these prompts still over-represent Western entities to some degree. Differential familiarity with these entities could elicit marginally different conversational patterns, which in turn might influence downstream properties of the dataset. Depending on the use case, users of this data may wish to consider supplementary sources to increase diversity and representation.

5. Conclusion

This paper presents the multi-dialect dataset of dialogues (MD3), which includes several thousand conversational information-sharing dialogues from English speakers in India, Nigeria, and United States. MD3 is distinguished by two key design decisions. First, the focus on information-sharing tasks makes it possible to define the *intent* of each dialogue. In future work we plan to test the accuracy and dialect robustness of systems for recovering this intent from the text of the conversation. A second distinction is the focus on nation-level dialects of global English. Much prior work on robustness in speech recognition has focused on what Wassink et al. call “ethnicity-related dialects” [30], such as African-American English. Our view is that global English is relatively understudied from a robustness perspective, and we hope that this dataset draws attention to this pervasive form of language variation. That said, the MD3 elicitation methodology could be directly applied to other classes of dialects, and we view this too as an interesting possibility for future work.

6. References

- [1] V. Pratap, A. Sriram, P. Tomasello, A. Hannun, V. Liptchinsky, G. Synnaeve, and R. Collobert, “Massively multilingual ASR: 50 languages, 1 model, 1 billion parameters,” *arXiv preprint arXiv:2007.03001*, 2020.
- [2] L. M. Arslan and J. H. Hansen, “Language accent classification in American English,” *Speech Communication*, vol. 18, no. 4, pp. 353–367, 1996.
- [3] A. Canavan and G. Zipperlen, “CALLFRIEND American English-Non-Southern Dialect,” Linguistic Data Consortium, Tech. Rep. LDC96S46, 1996.
- [4] W. Labov, *The social stratification of English in New York City*. Cambridge University Press, 2006.
- [5] A. Le, A. Martin, H. Hadfield, J. de Villiers, J.-P. Hosom, and J. van Santen, “2005 NIST Language Recognition Evaluation,” Linguistic Data Consortium, Tech. Rep. LDC2008S05, 2008.
- [6] S. Wray and A. Ali, “Crowdsource a little to label a lot: Labeling a speech corpus of dialectal Arabic,” in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [7] B. Schuppler, M. Hagmüller, J. A. Morales-Cordovilla, and H. Pessentheiner, “GRASS: the Graz corpus of read and spontaneous speech,” in *LREC*, 2014, pp. 1465–1470.
- [8] T. Samardzic, Y. Scherrer, and E. Glaser, “ArchiMob - a corpus of spoken Swiss German,” in *Proceedings of the tenth international conference on language resources and evaluation (LREC 2016)*. European Language Resources Association (ELRA), 2016.
- [9] M. A. Zissman, T. P. Gleason, D. M. Rekart, and B. L. Losiewicz, “Automatic dialect identification of extemporaneous conversational, Latin American Spanish speech,” in *1996 IEEE International Conference on Acoustics, Speech, and Signal Processing Conference Proceedings*, vol. 2. IEEE, 1996, pp. 777–780.
- [10] D. Demszky, D. Sharma, J. Clark, V. Prabhakaran, and J. Eisenstein, “Learning to recognize dialect features,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, Jun. 2021, pp. 2315–2338. [Online]. Available: <https://aclanthology.org/2021.naacl-main.184>
- [11] T. Masis, A. Neal, L. Green, and B. O’Connor, “Corpus-guided contrast sets for morphosyntactic feature detection in low-resource English varieties,” in *Proceedings of the first workshop on NLP applications to field linguistics*. Gyeongju, Republic of Korea: International Conference on Computational Linguistics, Oct. 2022, pp. 11–25. [Online]. Available: <https://aclanthology.org/2022.fieldmatters-1.2>
- [12] H. K. Craig and J. A. Washington, “An assessment battery for identifying language impairments in African American children,” *Journal of Speech, Language, and Hearing Research*, vol. 43, no. 2, pp. 366–379, 2000.
- [13] A. Koenecke, A. Nam, E. Lake, J. Nudell, M. Quartey, Z. Mengesha, C. Touns, J. R. Rickford, D. Jurafsky, and S. Goel, “Racial disparities in automated speech recognition,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 14, pp. 7684–7689, 2020.
- [14] A. Johnson, K. Everson, V. Ravi, A. Gladney, M. Ostendorf, and A. Alwan, “Automatic dialect density estimation for African American English,” in *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, H. Ko and J. H. L. Hansen, Eds. ISCA, 2022, pp. 1283–1287. [Online]. Available: <https://doi.org/10.21437/Interspeech.2022-796>
- [15] S. Greenbaum and G. Nelson, “The International Corpus of English (ICE) project,” *World Englishes*, vol. 15, no. 1, pp. 3–15, 1996.
- [16] T. Kendall and C. Farrington, “The Corpus of Regional African American Language,” The Online Resources for African American Language Project, Eugene, Oregon, Tech. Rep. Version 2021.07, 2021.
- [17] P. Budzianowski, T.-H. Wen, B.-H. Tseng, I. Casanueva, S. Ultes, O. Ramadan, and M. Gašić, “MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics, 2018, pp. 5016–5026. [Online]. Available: <https://aclanthology.org/D18-1547>
- [18] R. Tatman and C. Kasten, “Effects of Talker Dialect, Gender & Race on Accuracy of Bing Speech and YouTube Automatic Captions,” in *Proc. Interspeech 2017*, 2017, pp. 934–938.
- [19] P. Kaur, K. Sikka, W. Wang, S. Belongie, and A. Divakaran, “FoodX-251: a dataset for fine-grained food classification,” *arXiv preprint arXiv:1907.06167*, 2019.
- [20] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie, “The Caltech-UCSD Birds-200-2011 dataset,” California Institute of Technology, Tech. Rep. CNS-TR-2010-001, 2011.
- [21] A. Khosla, N. Jayadevaprakash, B. Yao, and F.-F. Li, “Novel dataset for fine-grained image categorization: Stanford dogs,” in *Proc. CVPR workshop on fine-grained visual categorization (FGVC)*, vol. 2, no. 1. Citeseer, 2011.
- [22] T. Bernard, “Tabouid: a Wikipedia-based word guessing game,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Online: Association for Computational Linguistics, Jul. 2020, pp. 24–29. [Online]. Available: <https://aclanthology.org/2020.acl-demos.4>
- [23] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” *arXiv preprint arXiv:2212.04356*, 2022.
- [24] U. Gut, “English in West Africa,” in *The Oxford Handbook of World Englishes*, D. S. M. Filppula, J. Klemola, Ed. Oxford: Oxford University Press, 2017, pp. 491–507.
- [25] P. Sailaja, *Indian English*. Edinburgh: Edinburgh University Press, 2009.
- [26] D. Sharma, “Typological diversity in New Englishes,” *English World-Wide*, vol. 30, no. 2, pp. 170–195, 2009.
- [27] U. Gut and R. Fuchs, “Progressive aspect in Nigerian English,” *Journal of English Linguistics*, vol. 41, no. 3, pp. 243–267, 2013.
- [28] C. Lange, “Focus marking in Indian English,” *English World-Wide*, vol. 28, no. 1, pp. 89–118, 2007.
- [29] N. Taguchi, “A comparative analysis of discourse markers in English conversational registers,” *Issues in Applied Linguistics*, vol. 13, no. 2, pp. 170–195, 2002.
- [30] A. B. Wassink, C. Gansen, and I. Bartholomew, “Uneven success: automatic speech recognition and ethnicity-related dialects,” *Speech Communication*, vol. 140, pp. 50–70, 2022.