

Check Mate: Prioritizing User Generated Multi-Media Content for Fact-Checking

Tarunima Prabhakar, Anushree Gupta, Kruttika Nadig, Denny George

Tattle Civic Technologies

tarunima@tattle.co.in, anushree.gupta93@gmail.com, kruttika17@gmail.com, denny@tattle.co.in

Abstract

Volume of content and misinformation on social media is rapidly increasing. There is a need for systems that can support fact checkers by prioritizing content that needs to be fact checked. Prior research on prioritizing content for fact-checking has focused on news media articles, predominantly in English language. Increasingly, misinformation is found in user-generated content. In this paper we present a novel dataset that can be used to prioritize check-worthy posts from multi-media content in Hindi. It is unique in its 1) focus on user generated content, 2) language and 3) accommodation of multi-modality in social media posts. In addition, we also provide metadata for each post such as number of shares and likes of the post on ShareChat, a popular Indian social media platform, that allows for correlative analysis around virality and misinformation. The data is accessible on Zenodo (<https://zenodo.org/record/4032629>) under Creative Commons Attribution License (CC BY 4.0).

Introduction

Misinformation has emerged as an unfortunate by-product of growing social media usage. Over the last few years, rapid Internet penetration and mobile phone adoption in developing countries such as India has led to a rapid increase in social media users. Many such users use regional Indian languages as their primary language of communication. The resultant misinformation has had life threatening consequences (Banaji et al. 2019). On chat apps such as WhatsApp, user-generated multi-media content has emerged as an important category of misinformation. Figure 1 shows an example of a heavily circulated post on WhatsApp, where a video is recirculated with a misleading text message. The video and text together constitute a post that needs to be verified.

Misinformation is contextual- the content and form of misinformation varies by topic of misinformation, platform of circulation, as well as the demographic targeted with that misinformation (Shu et al. 2017). The last few years have seen efforts towards characterizing and detecting ‘fake news’ on social media. There is a consistent need for datasets emerging from diverse contexts to better understand the phenomenon.

Fact-checking has emerged as an important defense mechanism against misinformation. Social media platforms rely on third party fact checkers to flag content on them. Fact checking groups also run tiplines¹ on WhatsApp through which they can source content that requires fact checking. The amount of content received by fact checkers however, significantly exceeds journalistic resources. Automation can support fact checkers by flagging content that is of higher priority for fact-checking. Extracting claims and assessing check-worthiness of posts is one way to help prioritise specific content for fact-checking (Arslan et al. 2020). In case of news articles and speeches, these claims appear as factual statements made by writers or speakers. User generated content, however, is often multi-modal. Videos and images are supported with additional text or media to give false context, make false connections or create misleading content (Wardle 2017). In Figure 1, the text message provides false context to an authentic video. In Figure 2, names and actions are attributed to people identified in images. In this specific case, the image of the individual in the top right is of a fashion designer and not of a politician, as claimed in the post.

Previous work has predominantly focused on claims detection for single media sources such as tweet texts, speeches and news articles (Arslan et al. 2020; Ma et al. 2018). Additionally, most research has focused on English language content. The rise in multi-modal user generated content makes it imperative to consider claims in multi-media posts. Furthermore, with regards to prioritizing multi-media content for fact checking, there is a need to deliberate on what is check-worthy in multi-media content. For example, selfie video claim lesser journalistic integrity than televised news coverage and might be of lower priority for fact checking.

In this paper, we present a unique dataset of multi-modal content in Hindi that can be used to train models for identifying check-worthy multi-modal content in the language. It is unique in its 1) focus on user generated content, 2) language and 3) claims detection in multi-modal content. In addition to the annotations, we also provide metadata for each post- such as time-stamp of creation; number of shares and likes; and the tags that accompany them on a popular Indian social

¹<https://techcrunch.com/2019/04/02/whatsapp-adds-a-tip-line-for-checking-fakes-in-india-ahead-of-elections/>

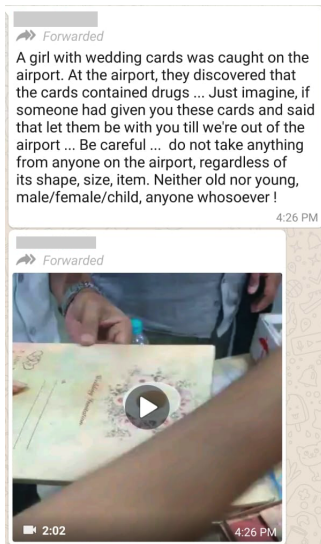


Figure 1: False Context in Multi-Modal Content



Figure 2: Text Embedded in Image

media platform, ShareChat. Table 1 summarizes the different media types in the dataset. Table 2 provides an overview of the fields in the dataset. In the subsequent section, we provide background on data collection and discuss related work. We then describe the process of annotation, the results from the annotations; and possible use cases and limitations.

Related Work

Background on Sharechat

The posts in this dataset were collected from an Indian social media platform called ShareChat. The platform has over 60 million monthly active users and is available in fourteen regional Indian languages (Banerjee 2020). It provides users the option to share other users' content on their profiles, as well as share content to other apps such as WhatsApp. (Agarwal et al. 2020) provide a detailed summary of linguistic and temporal patterns observed in the content on the

platform. For this dataset, we focus specifically on content in Hindi. With 500 million speakers, Hindi is the most commonly spoken Indian language (Indian-Census 2011).

Characterizing Check-Worthiness

Claim detection is a rich area of research (Aharoni et al. 2014; Levy et al. 2014). Given their importance to the American electoral process, US presidential debates have received significant attention for claim detection (Nakov et al. 2018; Patwari, Goldwasser, and Bagchi 2017; Arslan et al. 2020). Some research has focused on claims in user generated content on Twitter. This work however is still focused on textual claims on the platform (Ma et al. 2018; Lim, Lee, and Hsu 2017).

The work of (Zlatkova, Nakov, and Koychev 2019) is noteworthy for its focus on textual claims about images. In the field of misinformation detection, multi-media content has received more attention in the last few years. (Nakamura, Levy, and Wang 2020) provide a dataset of one million multi-modal posts from Reddit, classified into different types of misinformation. (Reis et al. 2020) provide a dataset of Brazilian and Indian election related images circulated on WhatsApp, that were found to be misinformation. (Gupta et al. 2013) analyze fake images on Twitter during Hurricane Sandy.

More broadly, literature describing and classifying misinformation highlights features that could help in prioritization of multi-media user-generated content for fact checking. (Shu et al. 2017) focus specifically on fake news articles which contain two major components- publisher and content. They classify news features using four attributes: source, headline, main text and image/video contained in the article. In their characterization of fake news, the image/video is considered as a supporting visual cue to the main news article. This is different from user generated content on chat apps, where the text can be supplementary context for an image or video. (Molina et al. 2019) propose an eight category typology for identifying 'fake news' which ranges from citizen journalism to satire and commentary. They further identify content attributes such as factuality, message quality, evidence, source and intention to characterize these eight categories.

Since a lot of misinformation research has focused on English content, we also highlight here work that focuses on non-English media. The dataset provided by (Reis et al. 2020) has images with content in Portuguese and Indian languages. (Abonizio et al. 2020) look at language features across English, Spanish and Portuguese for fake news detection.

Data Collection

To source user-generated multi-media content, we scraped content from ShareChat. On ShareChat, content is organized by 'buckets' corresponding to topic areas. Each bucket has multiple tags. Figure 3 shows the conceptual organization of content on the platform. We scraped content daily from the 'Health', 'Coronavirus' and 'Politics and News' buckets in Hindi from March 21 2020 to August 4 2020. We focused on

	Count
Total Number of Posts	2200
Annotated by Single Person	1800
Annotated by Two People	4000
Posts with primary media type: Image	1354
Posts with primary media type: Video	659
Posts with primary media type: Text	187

Table 1: Dataset Summary

Field Category	Field Names
Fields Describing Post (from platform)	bucket name, media type, caption, tag name, tag translation, text
Annotated Fields	verifiable claim, contains video, contains image, visible source, contains relevant meme
Metadata Fields (from platform)	number of external shares, likes, estimated timestamp of creation
Other Fields	Unique Post ID, Annotator Label

Table 2: Dataset Fields Overview

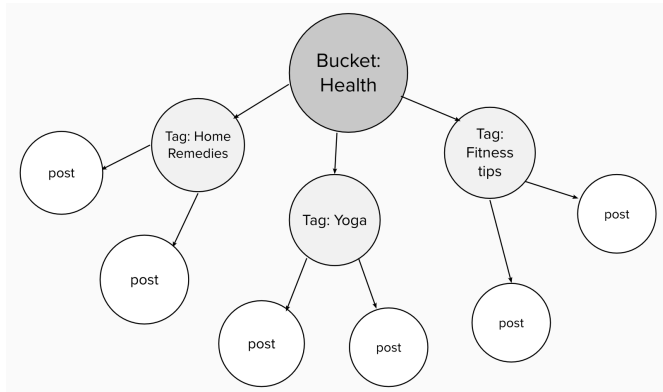


Figure 3: ShareChat Content Tree Structure

these buckets since politics and health have been observed to be the most common themes of misinformation in India (Chakrabarti, Stengel, and Solanki 2018). This resulted in over 200,000 unique posts. We randomly sampled from this bigger corpus to prepare the dataset presented here. The platform provides a number of additional data points such as the number of shares and likes, that provide some social context for a post. The platform also assigns a ‘media type’ to each post which could be text, image, video or link. Regardless of the media type, there is an option for users to provide captions for a post, which can contain large text descriptions. For our sampling, we ignore posts of media-type ‘link’ since these redirect to other websites and are therefore not within the purview of our focus on user generated content on the platform.

Qualitative Analysis

Open Coding

We use ‘Open coding’ methodology to identify features that can help prioritize multi-media content for fact checking. Open Coding is an analytic qualitative data analysis method to systematically analyze data (Mills, Durepos, and Wiebe 2010). In the first round, two researchers from the team loosely annotated and described 600 posts to familiarize themselves with the data. In the second phase, we described 300 posts in detail. We converged on sixteen non-mutually exclusive categories of multi-media posts that could be verified. These included posts that contained statistical claims (such as Covid fatality numbers); quotes attributed to noteworthy individuals; and newspaper clips and edited recordings from televised news reporting.

We also recognized user generated videos claiming to debunk misinformation as a separate category of content- the period of data collection coincided with Covid-19 and some posts were warnings about misinformation on the pandemic.

From these sixteen categories we converged on the following broad attributes of check-worthy content:

- The implied/explicit source of the content: Similar to (Shu et al. 2017; Molina et al. 2019) we find the *implied* source of the content to be an important attribute for prioritizing content for fact checking. We identified five common sources of data (see Table 3) that are ascribed differing levels of authenticity. Circulars and messages attributed to government sources for example are more likely to be believed than anonymous sources.
- The kind of verifiable claim: (Arslan et al. 2020) specifically characterize sentences from US presidential debates into three categories- non-factual statement, unimportant factual statement and check-worthy factual statement. Since this dataset is multi-modal, we found it useful to further qualify check-worthy factual statements.

We define three broad categories of check-worthy claims in multi-media content: 1. statistical/numerical claims. 2. world events or places about noteworthy individuals. 3. Other factual claims such as home remedies, cures and nutritional advice.

- For video content, the ‘intentionality’ of creation, as inferred from the video: There are multiple ways to categorize user generated video content. We broadly divide the content into two groups- videos where the person in the video is performing for the camera, and videos where the recording is incidental to the event. Given the ability to provoke, we find it useful to highlight violent incidents portrayed in videos as a third category. Notably, we do not consider an image montage to be a video, even if the media type as declared by the platform is ‘video’.
- Specifically for non digitally-manipulated images, the topic of images: While a lot of content on ShareChat is classified as ‘image’, not all of it is unidentified. A lot of ‘Good Morning Wishes’ memes use generic background images. In our framework, such images are not considered important for fact checking. We draw a distinction between digitally manipulated images and non-digitally manipulated images.
We realize that within non-digitally manipulated images, the topic of the image is an important feature that could help prioritize a content for fact checking. This however would require developing a content taxonomy of social media content. This is an important and complex research area and merits dedicated attention. We found this taxonomy to be beyond the ambit of this specific work. We loosely classified non-digitally manipulated images into those that contain human figures and those that don’t. We found it useful to flag images with human figures in them, since the context ascribed to the photograph can be verified.
- Culturally Relevant Memes: In the absence of a relevant taxonomy, we defined a broad binary category for ‘relevant memes’. For the purposes of this annotation, relevant memes are those that reference security forces, religious pride, national pride, political party representatives, reverential posts about a profession, election campaigning or current national news. While a more detailed taxonomy is needed, this broad descriptor of culturally important themes, when combined with other attributes such as source and verifiable claims, can become important in highlighting a post that needs to be fact-checked.

Table 3 summarizes these five attributes.

Closed Coding

Based on the above mentioned guidelines, the two researchers independently labeled posts in batches of fifty to refine these definitions. The framework developed was expanded into an annotation guideline with descriptions of each of the attributes to be labeled as well as sample annotations for thirteen posts. The annotation guide is provided with the dataset. The annotation guideline was opened to two additional members of the Tattle team, who were not

a part of the open coding stage. By virtue of their work, each member is familiar with the misinformation challenge in India, but not equally familiar with the nature of the content circulated. The four person annotation team was then split into two pairs. To estimate agreement across annotators, each person in the pair annotated 200 hundred posts common with the other member of the pair. Table 4 lists the agreement scores for the two pairs of annotators across the different attributes. The agreement for video type and image type attribute was substantial (greater than 0.6) for both pairs. Annotators strongly agreed on a post containing a verifiable claim as well. This can be used to train models for claim detection on multi-media posts. There was far weaker agreement on the specific kind of factual claim, when a post did contain one. Finally, on the implied or explicit source of content, one pair had moderate agreement while the second had weak agreement.

Annotation Procedure

After observing moderate to high agreement across four fields, each annotator proceeded to annotate 450 additional posts. This resulted in a total of 1800 posts uniquely annotated by one of the four annotators. We developed an in-house user interface for annotating specific posts². The interface is a Gatsby App built using React and Grommet UI library. Through the interface, annotators could load post related data from a JSON file and preview the media and metadata associated with it. Annotators could navigate from one post to another using ‘previous’ and ‘next’. On clicking ‘next’ annotated labels gets saved in the same JSON file. (See figure 6)

Attributes such as video style, image content and source that had mutually exclusive options were enforced as single option questions in the interface. The attribute about the kind of factual claim was listed as a check-box, since a post can contain multiple kinds of factual claims.

After all the annotations were completed, the JSON files were converted to CSV files for greater legibility and ease of sharing.

Dataset Description

Data Sample

In this dataset we share 400 posts that were annotated by two people each in the closed coding phase, as well as 1800 posts annotated by four separate annotators. In total, there are 2200 annotated posts shared in this dataset.

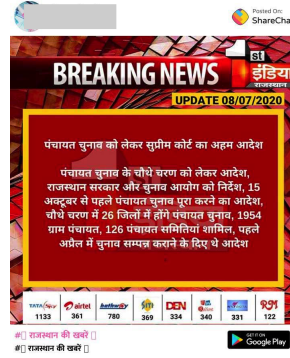
Before expanding on the annotations, we’ll describe the data fields scraped from ShareChat. The metadata about a post scraped from the platform includes:

- timestamp: time at which the post was created on the platform
- media type: can be image, video or text
- text: text content if media type of post is text.
- caption: text caption to images and videos.

²<http://annotation.test.tattle.co.in.s3-website.ap-south-1.amazonaws.com/>



(a)



(b)

Figure 4: (a) Post Attributed to Government Body (Niti Aayog) (b) A Screen Grab of TV News With Factual Claims



असम में ऑयल इंडिया के कुएं में आग : 2 फायर फाइटर की मौत, 1 खुलसा; कंपनी ने कहा- आग बुझाने में 4 हफ्ते लगेंगे

(a)



सुप: संघर्ष न्यून सुप
 डा. आंबेडकर के साथ भेदभाव और अन्याय आखिर किसने और कब किया था ? भेदभाव का कहानी किसने गढ़ा??
 डा. आंबेडकर के बारे में अक्सर कहा जाता है कि- उनके साथ भेदभाव और अन्याय किया जाता था?
 उनके ऊपर अध्ययन करने पर तो मुझे आज तक ऐसा कुछ नहीं मिला, जिससे पता चले कि किसी ने उनके साथ भेदभाव या अन्याय किया हो। यदि आय में से किसी के पास इस सम्बन्ध में कोई जानकारी हो तो बताने की कृपा करें।
 उनके बारे में पढ़ने पर पता चलता है कि - उनके साथ अगर जीवन में किसी ने अन्याय किया तो वो उनके पिता राजजी सकपाल ने किया था। जिसने 14 बच्चों का बाप होते हुए दूसरी शादी की और दूसरी बीवी के चक्कर में अपने बच्चों को छोड़ दिया। अपने पिता की उम्मीद और बीमारी के चलते 9 बच्चे अकाल मौत का शिकार हो गए।

(b)

Figure 5: (a) An Unidentified Image with Text Description About an Oil Well Accident (b) An Image with Text Description about Dr. Bhimrao Ambedkar

- external shares: number of times the post was shared from the platform to an external app such as WhatsApp
- number of likes: number of times the post was liked on ShareChat
- tags: the tags attached to the post by the user. See figure 3
- bucket name: the bucket under which this post was classified by the platform. See figure 3

Regardless of the media type attributed by the platform, users can add additional text for images and videos which is captured in the ‘caption’ field for each post.

Annotated Data

The fields developed in the closed coding stage were annotated as following:

- verifiable claim: a list of kinds of verifiable claims contained in the post. If no claim was found, the post is labeled as ‘No’. The claims can be those listed in Table 3. Since a post can have more than one kind of claim, the annotated field is a list of claims.
- contains video: if the post contains a video, it can be marked as ‘Person Performing for Camera’ or ‘Contains Violent Incident’. A video that doesn’t fall in these two categories is marked as other. A post that does not contain a video is marked as ‘No’.
- contains image: A post is said to contain an image if the image is **not** digitally manipulated. Within an image, a post can be marked as ‘containing human figures’ or ‘not containing human figures’. A post that does not contain an image is marked as ‘No’.
- visible source: Marked as one of the sources listed in Table 3

Attribute	Options
Implied or Explicit Source of Content(mutually-exclusive)	Print/TV Media: Newspaper clips or screen grabs from print or TV media sources. Examples of TV media are BBC, CNN, Aaj Tak (Indian); examples of print media: The Times of India, Dainik Jagran, Rajasthan Times. See Figure 4.b
	Digital Only News Outlets such as Huffington Post, Editor ji, myUpchar.
	Digital Content Providers such as Buzzfeed Goodful, or Instagram and Facebook pages.
	Government/Public Authority: Government logos or stamps visible in the post or content attributed to spokespersons, officials, civil servants/bureaucrats, other public office holders. OR Multilateral institutions and international public bodies such as WHO, ILO and World Bank. See Figure 4.a
	Anonymous/Social Media Users. See Figure 5.a
	Other
Type of Factual Claim (not mutually-exclusive)	Statistical or numerical claim. See Figure 4.a
	Descriptions of real world places, events or a noteworthy individual(s). See Figure 5.b
	Other factual claims
Intentionality Portrayed in the Video (mutually-exclusive)	Person is Performing for Camera
	Contains Violent Incident
	Other
Non Manipulated Images	Contains Humans. See Figure 5.b and Figure 4.1
	Does not Contain Humans. See Figure 5.a
Contains Relevant Memes?	Yes/No

Table 3: Annotation Guideline

	Video Type	Image Type	Source	No Factual Claim	Statistical Claim	Real Event, Place, Person	Other Factual Claim	Relevant Memes
Pair 1	0.74	0.68	0.5	0.76	0.65	0.30	0.54	0.70
Pair 2	0.85	0.68	0.35	0.83	0.33	0.31	0.50	0.61

Table 4: Agreement

- contains relevant meme: a binary category which is marked yes if the post is a text or image (not video) based meme about a culturally relevant topic.

Figure 8 summarizes the attributes of the annotated posts (n=1800) that had high agreement in closed coding stage. Even though a majority of posts are labeled as images by the platform, only forty percent of the posts have non-manipulated images that can be fact checked (see Figure 7(b)). In this sample of data, for a majority of posts, the visible source of the post was the platform user (see Figure 7(a)). It must be clarified that the user could have picked up their content from other established sources such as online websites and newspapers. Those sources however were not visible in the post. Part of the challenge of fact-checking user generated content is identifying sources from which content emerges. The high prevalence of user generated content

without a recognized source, emphasizes the need for systems that prioritize this content for fact checking.

Dataset Structure

The Tattle CheckMate Database consists of two folders. The first folder contains annotations from the closed coding phase. The second folder contains annotations of 1800 posts done by single annotators. The dataset also contains a README file that describes the folder structure and fields in the datasets.


Ethical Considerations

Data Collection and Sharing

While content creators need to be logged into the platform to upload posts, anyone with access to the post URL can view it. We thus treat the posts as content in the public domain

Posts

previous 5f20ae871b2328c5bd9fc5a7 next 2/449



Key	Value
URL	https://sharechat.com/post/ZR6qYMR
Tag	भारत न्यूज़
Timestamp	2020-07-09 14:10:08
Caption	जब राज्य के मुख्यमंत्री ही अपनी प्रतिभा का मजाक बनाए, फिर जवाब तो बदला है भैया... Prashant Kishore Replies Hard to Nitish Kumar नरुन सुनं और शेयर करें, ताकि बिहार का हर युवा इसे याद रखे... #PrashantKishor #TeamPK #BaatBiharKi
Text	

Post contains a verifiable claim

No

Yes, Statistical Claim

Yes, Descriptions of a real world events/place about noteworthy individuals.

Yes, other

Post contains unidentified videos

No

Yes, person is performing for camera

Yes, contains violent incident

Other

Post contains images

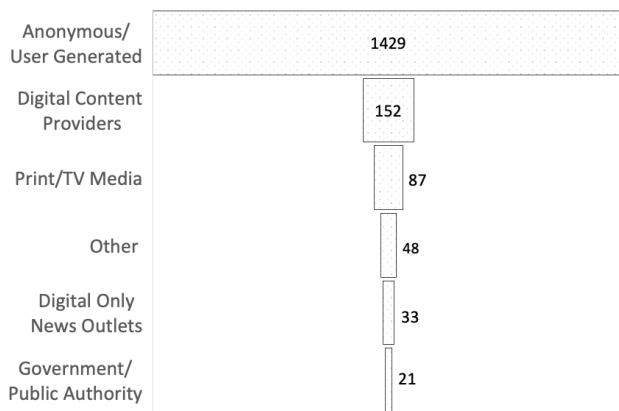
No, does not contain images

Contains human

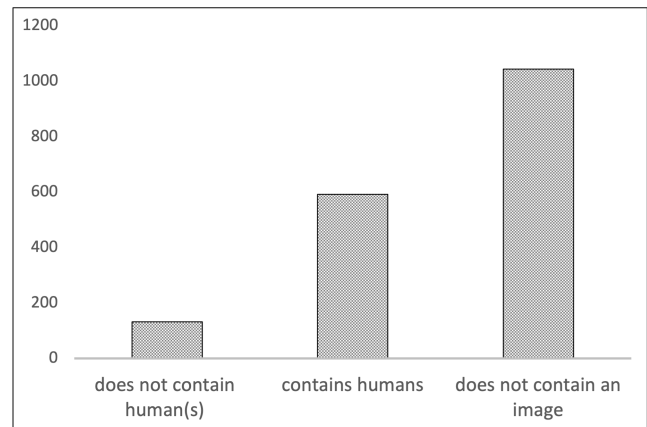
Does not contain human(s)

The visible source is

Figure 6: Annotation Interface



(a)



(b)

Figure 7: (a) Visible Source (b) Type of Non-Manipulated Images

that can be scraped and researched. We recognize that this content may contain images and videos of platform users and that opening the data can introduce some risk of identification of people in these posts. We have however mitigated this risk by removing the original source URL and sharing only the post. We also recognize that the violent imagery in some of the posts could be disturbing for users of this data. While labels such as those for violent videos could serve as an alert or trigger warning for some posts, we recognize that this is not comprehensive. Disturbing imagery is however unavoidable in the problem space that this dataset aims to intervene in. It is, in fact, one aspect of social media use that we hope this dataset can inform.


FAIR Requirements

This dataset adheres to FAIR principles (Findable, Accessible, Interoperable and Re-usable). The dataset is Findable and Accessible through Zenodo on


<https://zenodo.org/record/4032629> and licensed under Creative Commons Attribution License (CC BY 4.0). The data is shared as a set of CSV files with the referenced media items in the folder. Thus the data is reusable and inter-operable.

Conclusion

To the best of our knowledge, this dataset is a first of its kind in focusing on claims detection and prioritization of multi-media content for fact checking. This dataset provides a starting framework for features in different mediums that can help in flagging content. Future work can expand on this framework for more granular features across the different modalities. This dataset is also novel in its cultural and linguistic context and can be used for a range of open research problems:



(a)



(b)

Figure 8: (a) Type of Video (b) Proportion of Posts Containing Verifiable Claims

Training Models to Prioritize Multi-modal Content for Fact-Checking/Verification

This dataset can be used for building models that can highlight check-worthy multi-modal content, especially in Hindi. This data can also supplement other datasets on misinformation. (Reis et al. 2020) summarize datasets that are focused on the automation of fake news detection.

Understanding Content Virality

This dataset provides content features as well as some social context such as the number of shares, likes and tags. It can be used to better understand the relationship between the content and style of post and its popularity.

This work provides a framework for working with multi-modal content, but it also highlights a number of open research questions such as the need to build a taxonomy of image based content, and a more detailed typology of verifiable claims that are relevant for prioritizing of multi-modal content for fact checking. Future research can expand on this work to support such questions with other datasets.

Acknowledgements

We thank Dr. Sandeep Avula and Dr. Swair Shah for their continuous feedback in creation of this dataset. This work was partially supported by the ‘The Ethics and Governance of the Artificial Intelligence Fund’ of the AI Ethics Initiative.

References

Abonizio, H. Q.; de Moraes, J. I.; Tavares, G. M.; and Barbon Junior, S. 2020. Language-independent fake news detection: English, portuguese, and spanish mutual features. *Future Internet* 12(5):87.

Agarwal, P.; Garimella, K.; Joglekar, S.; Sastry, N.; and Tyson, G. 2020. Characterising user content on a multi-lingual social network. *Proceedings of the International AAAI Conference on Web and Social Media* 14(1):2–11.

Aharoni, E.; Polnarov, A.; Lavee, T.; Hershovich, D.; Levy, R.; Rinott, R.; Gutfreund, D.; and Slonim, N. 2014. A

benchmark dataset for automatic detection of claims and evidence in the context of controversial topics. In *Proceedings of the First Workshop on Argumentation Mining*, 64–68. Baltimore, Maryland: Association for Computational Linguistics.

Arslan, F.; Hassan, N.; Li, C.; and Tremayne, M. 2020. A benchmark dataset of check-worthy factual claims. *Proceedings of the International AAAI Conference on Web and Social Media* 14(1):821–829.

Banaji, S.; Bhat, R.; Agarwal, A.; Passanha, N.; and Sadhana Pravin, M. 2019. Whatsapp vigilantes: An exploration of citizen reception and construction of whatsapp messages’ triggering mob violence in india. *London School of Economics* <https://www.lse.ac.uk/media-and-communications/research/research-projects/whatsapp-vigilantes>. Accessed: 2020-12-09.

Banerjee, P. 2020. Sharechat getting half a million new users an hour since ban on 59 chinese apps. *Mint*. <https://www.livemint.com/news/india/sharechat-getting-half-a-million-new-users-an-hour-since-ban-on-59-chinese-apps-11593600168690.html>. Accessed: 2020-12-09.

Chakrabarti, S.; Stengel, L.; and Solanki, S. 2018. Duty, identity, credibility- fake news and the ordinary citizen in india. *BBC*. <http://downloads.bbc.co.uk/mediacentre/duty-identity-credibility.pdf>. Accessed: 2020-12-09.

Gupta, A.; Lamba, H.; Kumaraguru, P.; and Joshi, A. 2013. Faking sandy: Characterizing and identifying fake images on twitter during hurricane sandy. In *Proceedings of the 22nd International Conference on World Wide Web, WWW ’13 Companion*, 729–736. New York, NY, USA: Association for Computing Machinery.

Indian-Census. 2011. Population by bilingualism and trilingualism. Office of the Registrar General and Census Commissioner, Government of India. <https://www.censusindia.gov.in/2011census/C-17.html>. Accessed: 2020-12-09.

Levy, R.; Bilu, Y.; Hershovich, D.; Aharoni, E.; and Slonim, N. 2014. Context dependent claim detection. In

Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers, 1489–1500. Dublin, Ireland: Dublin City University and Association for Computational Linguistics.

Lim, W. Y.; Lee, M. L.; and Hsu, W. 2017. Ifact: An interactive framework to assess claims from tweets. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM '17*, 787–796. New York, NY, USA: Association for Computing Machinery.

Ma, W.; Chao, W.; Luo, Z.; and Jiang, X. 2018. Claim retrieval in twitter. In Hacid, H.; Cellary, W.; Wang, H.; Paik, H.-Y.; and Zhou, R., eds., *Web Information Systems Engineering – WISE 2018*, 297–307. Cham: Springer International Publishing.

Mills, A. J.; Durepos, G.; and Wiebe, E. 2010. Encyclopedia of case study research. *SAGE Publications, Inc.*

Molina, M. D.; Sundar, S. S.; Le, T.; and Lee, D. 2019. “fake news” is not simply false information: A concept explication and taxonomy of online content. *American Behavioral Scientist* 0002764219878224.

Nakamura, K.; Levy, S.; and Wang, W. Y. 2020. Fakeddit: A new multimodal benchmark dataset for fine-grained fake news detection. In *Proceedings of the 12th Language Resources and Evaluation Conference*, 6149–6157. Marseille, France: European Language Resources Association.

Nakov, P.; Barrón-Cedeño, A.; Elsayed, T.; Suwaileh, R.; Márquez, L.; Zaghoulani, W.; Atanasova, P.; Kyuchukov, S.; and Da San Martino, G. 2018. Overview of the clef-2018 checkthat! lab on automatic identification and verification of political claims. In Bellot, P.; Trabelsi, C.; Mothe, J.; Murtagh, F.; Nie, J. Y.; Soulier, L.; SanJuan, E.; Cappelato, L.; and Ferro, N., eds., *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, 372–387. Cham: Springer International Publishing.

Patwari, A.; Goldwasser, D.; and Bagchi, S. 2017. A multi-classifier system for detecting check-worthy statements in political debates. In *Conference on Information and Knowledge Management (CIKM)*.

Reis, J. C. S.; Melo, P.; Garimella, K.; Almeida, J. M.; Eckles, D.; and Benevenuto, F. 2020. A dataset of fact-checked images shared on whatsapp during the brazilian and indian elections. *Proceedings of the International AAAI Conference on Web and Social Media* 14(1):903–908.

Shu, K.; Sliva, A.; Wang, S.; Tang, J.; and Liu, H. 2017. Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.* 19(1):22–36.

Wardle, C. 2017. Fake news. it’s complicated. *First Draft*. <https://firstdraftnews.org/latest/fake-news-complicated/>. Accessed: 2020-12-09.

Zlatkova, D.; Nakov, P.; and Koychev, I. 2019. Fact-checking meets fauxtography: Verifying claims about images. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2099–2108. Hong Kong, China: Association for Computational Linguistics.