

YouTube AV 50K: An Annotated Corpus for Comments in Autonomous Vehicles

Tao Li, Lei Lin, Minsoo Choi, Kaiming Fu, Siyuan Gong, Jian Wang
Purdue University

{taoli, lin954, choi502, fu241, gong131, wang2084}@purdue.edu

Abstract—With one billion monthly viewers, and millions of users discussing and sharing opinions, comments below YouTube videos are rich sources of data for opinion mining and sentiment analysis. We introduce the YouTube AV 50K dataset, a freely-available collections of more than 50,000 YouTube comments and metadata below autonomous vehicle (AV)-related videos. We describe its creation process, its content and data format, and discuss its possible usages. Especially, we do a case study of the first self-driving car fatality to evaluate the dataset, and show how we can use this dataset to better understand public attitudes toward self-driving cars and public reactions to the accident. Future developments of the dataset are also discussed.

I. INTRODUCTION

Social media has become prevalent and important for social networking and opinion sharing in recent years [1]. By changing the way we perceive and interact with the world, social media has changed our lives profoundly [2], [3]. With millions of posts and replies uploaded every day on social media such as Facebook, Twitters and YouTube, it is an abundant and informative data source of public opinions; thus, it has attracted lots of attention from both academia and industry to understand people and society [4]–[6]. Most previous text mining-based social media analysis focused on Twitter and Facebook [7]. YouTube, generally considered as a video platform, the values of its text comments below videos have long been underestimated. Being the second most popular website in the world [8] and having 1.9 billion active users [9], YouTube is an attractive source of research in social media analysis with immense potentials.

Recent developments in autonomous vehicle technology have helped bring self-driving vehicles to the forefront of public interest [10]. Autonomous vehicles, particularly after the first fatal crash of self-driving cars recently happened in Arizona [11], have become a popular topic in social media. In order to investigate users’ acceptance, safety concerns, and willingness to purchase of autonomous vehicles, extensive studies have been conducted [12]–[14]. However, these research rely heavily on surveys, which have disadvantages that (i) securing a high response rate is hard; (ii) uncertainty over the validity of the data and sampling issues; and (iii) concerns surrounding the design, implementation, and evaluation of the survey [15]–[18]. Latest techniques in opinion mining and sentiment analysis from social media data [19]–[23] offer new possibilities to overcome disadvantages of traditional surveys, and the YouTube AV 50K dataset is introduced under this background.

This paper is organized as follows. In Section II we introduce our motivation of building this dataset as well as the data creation process, annotation methods, data formats, and basic statistics of the dataset. We propose possible usages of the dataset with examples in Section III, including metadata analysis, text visualization, sentimental analysis, recommendation systems, and text regression. Related works are presented in Section IV. Finally, Section V describes our visions and future goals of this project.

II. THE DATASET

A. Motivation

The YouTube Autonomous Vehicle (AV) 50K dataset is our attempt to help researchers in both natural language processing (NLP) community and autonomous vehicle community by providing a large-scale annotated corpus in the autonomous vehicle domain specifically. The dataset contains comments that are legally available to the public below a comprehensive list of autonomous vehicle-related vedieos on YouTube. Main purposes of the dataset includes: (i) encourage researchers on theories and algorithms to connect their works with real-world data; (ii) provide a reference dataset and benchmarks for evaluating research; and (iii) help new researchers quickly get started in natural langauge processing, data analysis, and autonomous vehicles.

B. Corpus

The core of the corpus is from the YouTube Data APIs¹. The APIs can be used to search for videos matching specific search terms, topics, locations, publication dates, and much more. Other helpful features for social media analysis include metadata of video, comment, playlist, author, and channel. For example, original text, like count, publish date, author name, and viewer rating can be found in each comment representation. Basically, any public available resources on YouTube can be downloaded using the APIs. Tutorials and more information can be found in the API document².

C. Annotation

Sentiment score is a numerical representation of the sentiment polarity, the degree of negative, neutral, or positive, for a piece of text. We used the Natural Language API³ for

¹<https://developers.google.com/youtube/v3/>

²<https://developers.google.com/youtube/v3/docs/>

³<https://cloud.google.com/natural-language/>

Item	#
Channels	50
Data size (MB)	119
Videos	632
Annotated comments	19,126
Unique authors	21,005
First level comments	21,498
Total comments	30,456
Word count	796,371

TABLE I
BASIC STATISTICS OF YOUTUBE 50K.

```

{
  "kind": "youtube#comment",
  "etag": etag,
  "id": string,
  "snippet": {
    "authorDisplayName": string,
    "authorProfileImageUrl": string,
    "authorChannelUrl": string,
    "authorChannelId": {
      "value": string
    },
    "channelId": string,
    "videoId": string,
    "textDisplay": string,
    "textOriginal": string,
    "parentId": string,
    "canRate": boolean,
    "viewerRating": string,
    "likeCount": unsigned integer,
    "moderationStatus": string,
    "publishedAt": datetime,
    "updatedAt": datetime
  }
}

```

Fig. 1. JSON structure shows the format of a comments resource.

sentimental analysis. We consider the sentiment score given by the API as a benchmark and will discuss different models for sentimental analysis in our upcoming paper. Evaluation of the annotation will be shown in the following Section III.

D. Content

The YouTube AV 50K contains comments and metadata for all public self-driving-related videos on YouTube. Basic statistics is shown in Table I (the numbers might change as the dataset is constantly updating).


The data is stored using JSON format [24] described by YouTube Data API to efficiently handle the heterogeneous types of information such as text, user id, publish date, like count, etc. Each comment is described by a JSON object⁴, whose structure is shown in Figure 1.

The YouTube AV 50K website⁵ is a core component of the corpus. It contains tutorials, code samples, API documents, an issue section, and the pointers to the actual data, hosted by Goodata Foundation⁶.


⁴<https://developers.google.com/youtube/v3/docs/comments>

⁵<https://youtubeav50k.goodata.org>

⁶<https://goodata.org>



(a) number of comments by hour



(b) number of comments by weekday

Fig. 2. Distribution of comment publish time.

III. PROPOSED USAGES

A wide range of social media analysis tasks can be performed and measured on the YouTube AV 50K dataset. In this section, we propose some possible usages of the dataset, and illustrate and evaluate the power of the dataset.

A. Metadata Analysis

A basic usage of the dataset is to analyze the metadata. This task is already very interesting, because trending topics, geographic information, and even patterns of human behaviors can be discovered from the metadata of a large volume of comments and videos. Statistical descriptions and visualisations can also be done for related research. An example of analyzing human comment behaviors is given in Figure 2, which illustrates a clear pattern that YouTube users are most active at around 8pm and least active in the early morning. Surprisingly, the figure shows that YouTube users leave more comments in average on every weekday rather than weekends. Another example usage is to provide channel rankings using the video metadata from the dataset. We provide a ranking of top 20 English news channels in YouTube by their monthly views from May 9, 2018 to June 9, 2018 in Table II. The

#	Channel	Monthly Views	Subscribers
1	Inside Edition	395025524	3129845
2	CNN	155319328	3794747
3	ABC News	144872445	3850158
4	Barcroft TV	129030102	4627820
5	MSNBC	74032667	963021
6	Fox News	60536948	1242054
7	Crime Watch Daily	53485050	877800
8	BBC News	45344148	2619048
9	The Young Turks	45159957	3899180
10	Guardian News	42598679	136110
11	RT	40166881	2632872
12	TODAY	40073086	758845
13	The Royal Family Channel	38726120	314798
14	USA TODAY	38444609	634366
15	VICE News	36201734	3048964
16	TomoNews US	34828513	2230927
17	CBS News	34513896	752014
18	DramaAlert	34501164	3888026
19	CBS This Morning	34236213	390538
20	Vox	32228568	4166051

TABLE II

TOP 20 YOUTUBE ENGLISH NEWS CHANNELS RANKED BY MONTHLY REVIEWS (FROM MAY 9TH, 2018 TO JUNE 9TH, 2018)

results are exactly the same compared to data from KEDOO⁷ downloaded by the *YoutubeStat.py* module [25].

B. Text Visualization

Word cloud is a straightforward and appealing visualization method for text. It has served as a starting point for many studies in text mining and opinion mining [26]. Our dataset, which has large amount of text data about people’s attitudes toward autonomous driving, is an ideal playground for word cloud generation as well as research in sentimental analysis and opinion mining. Figure 3 provides word clouds of the dataset generated by using a wordcloud library in Python⁸. On March 19, 2018, a self-driving Uber vehicle killed a pedestrian in Arizona and this incident is believed to be the first fatality associated with self-driving technology [11]. Figure 3a is generated from comments one month before the incident while Figure 3b shows a word cloud of comments one month after. We can see a clear divergence of topics in these two figures where the happening of the incident was the watershed.

C. Sentimental Analysis

Finding out what other people think has always been an important part of information collection and decision making. Gathering opinions of hot, and probably controversial topics in the social media has aroused tremendous attention in the research community. In recent years, autonomous driving technologies have achieved great progresses and gradually become such a topic - popular and controversial. Self-driving cars have been put through millions of miles of road tests, and some believe that the technology has the potential to be safer than human drivers [27]. However, with the recent fatal incident, public concerns of safety are clearly increasing according to the time series boxplot in Figure 4, which also

shows an interesting process of destroying and rebuilding public trusts over time, illustrated by the boxes that went up to the peak in February, then dropped dramatically in March because of the incident, and finally went up again recently.

Sentiment scores used in the boxplot are floating-point numbers representing sentiment polarities ranging from -1 to 1, where -1 is extreme negative, 0 is neutral, and 1 is extreme positive. The scores are already given in the dataset and were generated by the sentiment analysis feature of the Google Cloud Natural Language API⁹. It is natural that some might argue the validity of the API. We admit that the API is not perfect according to our experiments, and will provide a detailed comparison of results from this API and other state-of-the-art models in our upcoming paper. For now, we just introduce a model and leave other concerns to our next paper. The model is

$$\Phi(\Theta) = \phi(\Theta) + \epsilon_\phi \quad (1)$$

where Θ is an embedding of text, $\phi(\cdot)$ is a prediction model, $\Phi(\cdot)$ is the real model to be found, and ϵ_ϕ is independent to Θ but dependent to ϕ . The ϵ_ϕ of the model given by the API is considered to be white noise.

D. Recommendation Systems

With the expectation of immense commercial values [28], recommendation systems for advertisement have been studied extensively, and two main paradigms have emerged: content-based recommendation and collaborative recommendation [29]. Traditional content-based recommendation systems are Information Filtering (IF) systems that need proper techniques for representing the items and producing the user profile, and some strategies for comparing the user profile with the item representation, including content analyzer, profile learner, and filtering component [30]. The YouTube AV 50K dataset can be used either to generate sentimental scores mentioned in section III-C for user profile enrichments or to directly encode comment texts and metadata as inputs for deep neural networks to finally output ratings in recommendation systems.

E. Text Regression

Text regression is a method of predicting a real-world continuous quantity associated with the text’s meaning based on a piece of text and it was reported to significantly outperform past volatility in predicting financial indices [31]. Empirically, sales of products could be greatly influenced by online reviews. The case would be especially true for vehicle industry, because vehicles are generally more valuable than normal goods, and intuitively, people would be more careful in the selection processes and tend to look up more reviews. Considering these features, the YouTube AV 50K dataset is appealing for vehicle market as a source of sales prediction and risk measurement research. One of our upcoming papers would combine Bayesian deep learning with text regression to address this issue [32], [33].


⁷<https://www.kedoo.com/youtube/>

⁸https://github.com/amueller/word_cloud

⁹<https://cloud.google.com/natural-language/>



(a) one month before the crash



(b) one month after the crash

Fig. 3. Word cloud generated from YouTube comments below self-driving related videos before and after the first self-driving car fatality on March 19, 2018 in Tempe, Arizona.




Fig. 4. Boxplot of sentiment scores changes over time.

IV. RELATED WORKS

The YouTube AV 50K is the first dataset dedicated for social media analysis in autonomous vehicle-related fields. Other very large datasets for social media analysis includes:

- Amazon Fine Food Reviews¹⁰: 568,454 food reviews Amazon users left up to October 2012.
- Amazon Reviews¹¹: Stanford collection of 35 million amazon reviews.
- Disasters on social media¹²: 10,000 tweets with annotations whether the tweet referred to a disaster event.
- Reddit Comments¹³: every publicly available reddit comment as of July 2015. 1.7 billion comments.
- Twitter Cheng-Caverlee-Lee Scrape¹⁴: Tweets from September 2009 - January 2010, geolocated.

¹⁰<https://www.kaggle.com/snap/amazon-fine-food-reviews>

¹¹<https://snap.stanford.edu/data/web-Amazon.html>

¹²<https://www.figure-eight.com/data-for-everyone/>

¹³https://www.reddit.com/r/datasets/comments/3bxlg7/i_have_every_publicly_available_reddit_comment/

¹⁴https://archive.org/details/twitter_cikm_2010

- Classification of political social media¹⁵: Social media messages from politicians classified by content.
- Corporate messaging¹⁶: A data categorization job concerning what corporations actually talk about on social media.

Above is a selected list and not intended to be comprehensive. More related open datasets can be found in following Github repositories:

- [niderhoff/nlp-datasets](https://github.com/niderhoff/nlp-datasets)¹⁷
- [awesomedata/awesome-public-datasets](https://github.com/awesomedata/awesome-public-datasets)¹⁸.

V. FUTURE WORKS

We want to expand the impacts of our previous research in autonomous vehicles [32]–[35] and have better understanding of what people think of self-driving cars. We are especially interested in people’s opinions toward the transitions from level 0: no automation to level 5: full automation [36], [37]. Although YouTube AV 50K already provides rich sources of data annotated by state-of-the-art NLP techniques, we are aware of the shortages of the techniques; for example, doing aspect-level sentimental analysis and detecting sarcasm are still challenging. We would keep an eye on latest progresses in the NLP community and be open and ready to apply such techniques to our dataset.

Another improvement we would like to perform is to keep expanding the dataset as the autonomous driving field is growing rapidly and many exciting achievements are expected, which might greatly change the self-driving landscape as well as public opinions. Hopefully, such changes would mitigate people’s security concerns and bring out safer, smoother, and smarter transportation.

ACKNOWLEDGMENTS

We thank authors of [38]–[41] for helpful discussions, and the anonymous reviewers for useful feedbacks.

¹⁵<https://www.figure-eight.com/data-for-everyone/>

¹⁶<https://registry.opendata.aws/commoncrawl/>

¹⁷<https://github.com/niderhoff/nlp-datasets>

¹⁸<https://github.com/awesomedata/awesome-public-datasets>

REFERENCES

- [1] S. Asur and B. A. Huberman, "Predicting the future with social media," in *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology-Volume 01*. IEEE Computer Society, 2010, pp. 492–499.
- [2] L. Safko, *The social media bible: tactics, tools, and strategies for business success*. John Wiley & Sons, 2010.
- [3] W. He, S. Zha, and L. Li, "Social media competitive analysis and text mining: A case study in the pizza industry," *International Journal of Information Management*, vol. 33, no. 3, pp. 464–472, 2013.
- [4] A. Perrin, "Social media usage," *Pew research center*, pp. 52–68, 2015.
- [5] S. Boulianne, "Social media use and participation: A meta-analysis of current research," *Information, Communication & Society*, vol. 18, no. 5, pp. 524–538, 2015.
- [6] A. Ceron, L. Curini, S. M. Iacus, and G. Porro, "Every tweet counts? how sentiment analysis of social media can improve our knowledge of citizens political preferences with an application to italy and france," *New Media & Society*, vol. 16, no. 2, pp. 340–358, 2014.
- [7] D. Kulkarni and S. Rodd, "Extensive study of text based methods for opinion mining," in *2018 2nd International Conference on Inventive Systems and Control (ICISC)*. IEEE, 2018.
- [8] alexa.com. Alexa top 500 global sites. [Online]. Available: <https://www.alexa.com/topsites>
- [9] statista.com. Youtube - statistics & facts. [Online]. Available: <https://www.statista.com/topics/2019/youtube/>
- [10] B. Schoettle and M. Sivak, "A survey of public opinion about autonomous and self-driving vehicles in the us, the uk, and australia," 2014.
- [11] D. Wakabayashi, "Self-driving uber car kills pedestrian in arizona, where robots roam," *The New York Times*, 2018.
- [12] D. Howard and D. Dai, "Public perceptions of self-driving cars: The case of berkeley, california," in *Transportation Research Board 93rd Annual Meeting*, vol. 14, no. 4502, 2014.
- [13] M. Kyriakidis, R. Happee, and J. C. de Winter, "Public opinion on automated driving: Results of an international questionnaire among 5000 respondents," *Transportation research part F: traffic psychology and behaviour*, vol. 32, pp. 127–140, 2015.
- [14] B. Schoettle and M. Sivak, "Public opinion about self-driving vehicles in china, india, japan, the us, the uk, and australia," 2014.
- [15] K. B. Wright, "Researching internet-based populations: Advantages and disadvantages of online survey research, online questionnaire authoring software packages, and web survey services," *Journal of computer-mediated communication*, vol. 10, no. 3, p. JCMC1034, 2005.
- [16] K. Kelley, B. Clark, V. Brown, and J. Sitzia, "Good practice in the conduct and reporting of survey research," *International Journal for Quality in health care*, vol. 15, no. 3, pp. 261–266, 2003.
- [17] F. J. Fowler Jr, *Survey research methods*. Sage publications, 2013.
- [18] R. D. Fricker and M. Schonlau, "Advantages and disadvantages of internet research surveys: Evidence from the literature," *Field methods*, vol. 14, no. 4, pp. 347–367, 2002.
- [19] B. Pang, L. Lee *et al.*, "Opinion mining and sentiment analysis," *Foundations and Trends® in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.
- [20] A. Pak and P. Paroubek, "Twitter as a corpus for sentiment analysis and opinion mining," in *LREc*, vol. 10, no. 2010, 2010, pp. 1320–1326.
- [21] A. Esuli and F. Sebastiani, "Sentiwordnet: a high-coverage lexical resource for opinion mining," *Evaluation*, vol. 17, pp. 1–26, 2007.
- [22] S. Baccianella, A. Esuli, and F. Sebastiani, "Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining," in *Lrec*, vol. 10, no. 2010, 2010, pp. 2200–2204.
- [23] B. Liu, "Sentiment analysis and opinion mining," *Synthesis lectures on human language technologies*, vol. 5, no. 1, pp. 1–167, 2012.
- [24] D. Crockford, "The application/json media type for javascript object notation (json)," Tech. Rep., 2006.
- [25] T. Li and L. Lin, "Youtubestat.py: a python module to download youtube statistics and rankings," 2018. [Online]. Available: <https://github.com/Eroica-cpp/YouTube-Statistics>
- [26] F. Heimerl, S. Lohmann, S. Lange, and T. Ertl, "Word cloud explorer: Text analytics based on word clouds," in *System Sciences (HICSS), 2014 47th Hawaii International Conference on*. IEEE, 2014, pp. 1833–1842.
- [27] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, pp. 90–96, 2017.
- [28] B. Graham, "Social network e-commerce and advertisement tracking system," Jun. 28 2007, uS Patent App. 11/413,259.
- [29] M. Balabanović and Y. Shoham, "Fab: content-based, collaborative recommendation," *Communications of the ACM*, vol. 40, no. 3, pp. 66–72, 1997.
- [30] P. Lops, M. De Gemmis, and G. Semeraro, "Content-based recommender systems: State of the art and trends," in *Recommender systems handbook*. Springer, 2011, pp. 73–105.
- [31] S. Kogan, D. Levin, B. R. Routledge, J. S. Sagi, and N. A. Smith, "Predicting risk from financial reports with regression," in *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2009, pp. 272–280.
- [32] T. Li, "Modeling uncertainty in vehicle trajectory prediction in a mixed connected and autonomous vehicle environment using deep learning and kernel density estimation," *The Fourth Annual Symposium on Transportation Informatics*, 2018.
- [33] L. Lin, S. Gong, and T. Li, "Deep learning-based human-driven vehicle trajectory prediction and its application for platoon control of connected and autonomous vehicles," *The Autonomous Vehicles Symposium 2018*, 2018.
- [34] S. Gong, A. Zhou, J. Wang, T. Li, and S. Peeta, "Cooperative adaptive cruise control for a platoon of connected and autonomous vehicles considering dynamic information flow topology," *arXiv preprint arXiv:1807.02224*, 2018.
- [35] L. Lin, M. Ni, Q. He, J. Gao, and A. W. Sadek, "Modeling the impacts of inclement weather on freeway traffic speed: exploratory study with social media data," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2482, pp. 82–89, 2015.
- [36] B. M. Blumberg and T. A. Galyean, "Multi-level direction of autonomous creatures for real-time virtual environments," in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 1995, pp. 47–54.
- [37] J. Yan, R. Kowalczyk, J. Lin, M. B. Chhetri, S. K. Goh, and J. Zhang, "Autonomous service level agreement negotiation for service composition provision," *Future Generation Computer Systems*, vol. 23, no. 6, pp. 748–759, 2007.
- [38] X. Liu, T. Li, H. Peng, I. C. Ouyang, T. Kim, R. Wang, and G. Guo, "Mining semantic descriptions from data for beauty understanding," in *2019 IEEE Winter Conference on Applications of Computer Vision*, 2019.
- [39] T. Li, X. Liu, and S. Su, "Semi-supervised text regression with conditional generative adversarial networks," in *the 3rd International Workshop on Application of Big Data for Computational Social Science of the 2018 IEEE International Conference on Big Data*, 2018.
- [40] C. Wang, S. Gong, A. Zhou, T. Li, and S. Peeta, "Cooperative adaptive cruise control for connected autonomous vehicles by factoring communication-related constraints," in *the 23rd International Symposium on Transportation and Traffic Theory*, 2019.
- [41] J. Zhou, A. Kim, M. Ochoa, H. Jiang, and B. Ziaie, "An ultrasonically powered micropump for on-demand in-situ drug delivery," in *Micro Electro Mechanical Systems (MEMS), 2016 IEEE 29th International Conference on*. IEEE, 2016, pp. 349–352.