

The Potential of Using Vision Videos for CrowdRE: Video Comments as a Source of Feedback

Oliver Karras

TIB – Leibniz Information Centre for
Science and Technology
Hannover, Germany
Email: oliver.karras@tib.eu

Eklekta Kristo

Leibniz University Hannover
Software Engineering Group
Hannover, Germany

Jil Klünder

Leibniz University Hannover
Software Engineering Group
Hannover, Germany

Email: eklekta.kristo@stud.uni-hannover.de Email: jil.kluender@inf.uni-hannover.de

Abstract—Vision videos are established for soliciting feedback and stimulating discussions in requirements engineering (RE) practices such as focus groups. Different researchers motivated the transfer of these benefits into crowd-based RE (CrowdRE) by using vision videos on social media platforms. So far, however, little research explored the potential of using vision videos for CrowdRE in detail. In this paper, we analyze and assess this potential, in particular, focusing on video comments as a source of feedback. In a case study, we analyzed 4505 comments on a vision video from YouTube. We found that the video solicited 2770 comments from 2660 viewers in four days. This is more than 50% of all comments the video received in four years. Even though only a certain fraction of these comments are relevant to RE, the relevant comments address typical intentions and topics of user feedback, such as *feature request* or *problem report*. Besides the typical user feedback categories, we found more than 300 comments that address the topic *safety* which has not appeared in previous analyses of user feedback. In an automated analysis, we compared the performance of three machine learning algorithms on classifying the video comments. Despite certain differences, the algorithms classified the video comments well. Based on these findings, we conclude that the use of vision videos for CrowdRE has a large potential. Despite the preliminary nature of the case study, we are optimistic that vision videos can motivate stakeholders to actively participate in a crowd and solicit numerous of video comments as a valuable source of feedback.

Index Terms—Requirements engineering, vision video, crowd, video comment, feedback, classification

I. INTRODUCTION

In the recent years, several researchers [1]–[5] introduced vision videos into Requirements Engineering (RE) for soliciting feedback and stimulating discussions. According to Karras et al. [6, p. 2], “a vision video is a video that represents a vision or parts of it [(problem, solution, improvement)] for achieving a shared understanding among all parties involved by disclosing, discussing, and aligning their mental models of the future system”. These videos are used to elicit and validate an integrated view of the future system regarding new and diverging aspects among all parties involved [7].

In 2017, Schneider et al. [8] proposed the use of vision videos for CrowdRE to solicit feedback. Videos are a well-suited medium for CrowdRE since their communication richness and effectiveness [9]–[11], as well as their ease of use [12]–[14], facilitate dissemination, discourse, and feedback via social media platforms such as YouTube [5], [8], [15]. The use of

vision videos on social media platforms offers the potential to gather millions of views and solicit thousands of comments [5]. So far, Schneider and Bertolli [15] explored the production of vision videos for CrowdRE. They also mentioned the potentially beneficial use of social media platforms to disseminate vision videos for soliciting feedback. However, they explicitly state that the use of vision videos for CrowdRE is future work.

In this paper, we focus on this use of vision videos for CrowdRE. We are interested in the potential of using vision videos on social media platforms to solicit feedback in the form of video comments. We ask the following research question:

RQ: What is the potential of using vision videos on social media platforms to solicit feedback in the form of video comments for CrowdRE?

In a case study, we investigate this potential from three perspectives: (1) *quantity*, (2) *content*, and (3) *suitability*. All three perspectives are important to assess the potential of using vision videos for CrowdRE [16]. First, we must determine how well vision videos can motivate stakeholders to participate in the crowd, i.e., to watch the videos and write comments. Second, we must examine the content of the comments and their relevance to RE. Third, we must investigate how well the comments are suited for the current analysis techniques in CrowdRE. Based on these three perspectives, we contribute the following insights:

- (1) *Quantity*: Vision videos can solicit many of comments from various viewers in a short period of time.
- (2) *Content*: In addition to typical categories of user feedback, video comments also address the topic *safety*, which has not appeared in analyses of user feedback so far.
- (3) *Suitability*: Despite certain differences in the performance, machine learning algorithms can already classify video comments well in terms of content.

This paper is structured as follows: Section II discusses related work. While section III describes the case study design, section IV reports the findings. After presenting threats to validity in section V, the findings are discussed in section VI. Section VII concludes the paper.

II. BACKGROUND AND RELATED WORK

One of the four services of CrowdRE is to provide motivational instruments to engage stakeholders to actively participate in a crowd [17]. However, this motivation is a major challenge since simply providing a forum is not enough to solicit feedback and stimulate discussions [18]. While most CrowdRE approaches focus on the use of forums and app stores, few consider social media platforms [16], [19], [20]. Khan et al. [16] argue that social media platforms, with their entertaining and enjoyable activities, e.g., watching a video, provide a great opportunity to motivate stakeholders to participate in a crowd.

Vision videos are one specific kind of videos. These videos are frequently used to solicit feedback and stimulate discussion in RE practices such as workshops and focus groups [13], [21], [22]. Although different researchers [1]–[4] reported on the benefits of vision videos for soliciting feedback and stimulating discussions in group or individual meetings, they seldom provide details about the *quantity* and *content* of the feedback. Bennaceur et al. [1] used vision videos in group discussions to solicit as much feedback from the participants as possible. Darby et al. [2] produced a vision video to gather feedback on a visionary scenario from a group of stakeholders. While both papers present insights into the main themes of the overall feedback, they do not provide details on the *quantity* or *content* of the individual contributions in the feedback. Schneider et al. [3] as well as Busch et al. [4] used vision videos in individual meetings with single stakeholders to solicit feedback. Both papers examined the number of contributions in the feedback. Although both papers defined what constitutes a contribution, including ideas, questions, requirements and rationale, neither paper explored the *content* of the feedback in detail. This lack of detail about feedback may be a problem of its direct nature in meetings. In this context, it is far more difficult to grasp, capture, and classify feedback immediately [23].

Based on the benefits of vision videos for soliciting feedback and stimulating discussion in meetings, the question arises whether these benefits can also be transferred to social media platforms. On these platforms, (video) comments represent the main communication channel. These comments are similar to app reviews, forum posts, or tweets and thus could be another source of feedback for CrowdRE. Following this line of thought, Vistisen and Poulsen [5] manually examined 300 comments of a vision video in a case study. Their results indicate that the use of vision videos on social media platforms can solicit valuable feedback from viewers. However, a manual analysis is not feasible in the long term due to the potentially large number of comments. Thus, automated analyses are unavoidable for which video comments must be *suitable* [24], [25].

Several researchers already applied various automated methods on video comments to analyze and classify them for a wide variety of purposes. Below, we can only give a brief overview based on some approaches due to the large number of use cases in various research areas. There are two main use cases for analyzing video comments.

The first use case is *spam* detection. In this case, *spam* refers to any irrelevant, promotional, incoherent, offensive, or harmful comment [26]–[29]. These comments need to be identified and removed from the social media platforms. The second use case is *ham* detection. In this case, *ham* refers to any relevant comment that provides an opinion, question, concern, need or other relevant information from the viewers to the video producers [30]–[33]. In all cases, the researchers used and compared several algorithms such as *support vector machines*, *naive Bayes*, *logistic regression*, or *neural networks* to classify the video comments for their respective purpose. In accordance with each other, these researchers emphasize the potential of video comments and their automated analysis.

These insights motivated us to further explore the potential of using vision videos on social media platforms for CrowdRE. For this purpose, we carried out a case study.

III. CASE STUDY

The design and reporting follow the guideline for case study research in software engineering by Runeson and Höst [34].

A. The Case and Unit of Analysis

The case is represented by *The Boring Company*¹. This company has the mission to solve traffic issues by creating safe, fast-to-build, and cost-effective transportation, utility, and freight tunnels. On April 28, 2017, *The Boring Company* released an animated video² on their vision to solve the traffic problems in cities by building large networks of tunnels underneath them. This vision video is the unit of analysis. The video shows a car that uses a network of tunnels underneath a city to travel from one location to another (see Fig. 1).



Fig. 1: Snapshot from the vision video of *The Boring Company*

B. Data Analysis Procedure

In Fig. 2, we present the entire data analysis procedure consisting of four phases: Data collection, data cleaning, manual analysis, and automated analysis.

¹<https://www.boringcompany.com/>

²https://www.youtube.com/watch?v=u5V_VzRrSBI

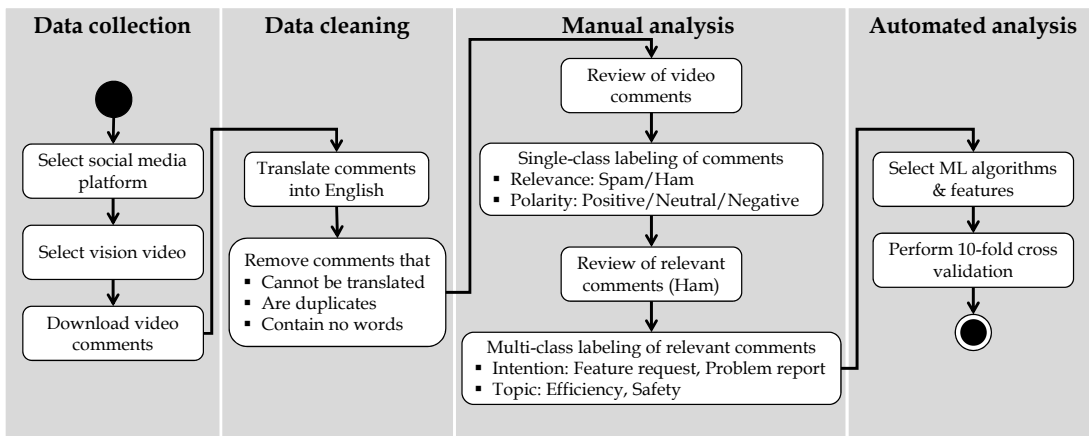


Fig. 2: Overview of the data analysis procedure

1) *Data Collection*: As social media platform, we selected YouTube since it is the most popular and widely used video-sharing platform where any user can upload, watch, comment, like, dislike, and discuss any videos [32], [35], [36]. Based on a manual search, we found the vision video shared by *The Boring Company* via its channel with 140000 subscribers. We chose this video due to its large number of comments. On October 13, 2020, we downloaded the data set of the video via the YouTube Data API³. TABLE I summarizes the statistics of the data set compared to our last access on the video exactly four years after its release. The data set contains 4505 video comments from 4302 different authors. In addition, there are more than 2000 replies that represent discussion threads below specific comments. In this paper, we only focus on the comments since these are directly related to the video. An analysis of the replies remains future work.

TABLE I: Overview of the statistics of the vision video

State	Views	Likes	Dislikes	Comments	Authors
Oct 13th, 2020 (download)	6.9 million	59621	4287	4505	4302
Apr 28th, 2021 (last access)	8.3 million	88457	4913	5184	4980

2) *Data Cleaning*: Although the majority of the comments is written in English, there are several comments written in different languages. We decided to translate these comments into English to keep as many comments in the data set as possible. For this purpose, we manually translated each comment using *Google Translate*⁴. For data cleaning, we removed all comments from the data set that could not be translated, were exact duplicates, or did not contain any words. The final data set contains 4400 comments for manual analysis.

3) *Manual Analysis*: The manual analysis served two purposes. First, this qualitative analysis allowed us to explore the potential of using vision videos for CrowdRE from the two

perspectives *quantity* and *content*. Second, we had to create a labeled data set for the automated analysis to investigate the third perspective *suitability*. We started with a manual review of the video comments to get an impression of their contents.

Based on the insights, we performed two classifications. For each classification, we provided the coders with a guide sheet that describes the coding task, defines the categories, and lists examples to reduce inconsistencies and increase the quality of manual labeling [37]. In the first classification, we distinguished between comments that are relevant (*ham*) to RE and those that are irrelevant (*spam*) [38]. For this purpose, two coders with a computer science background classified all 4400 video comments regarding their relevance and, in addition, their polarity. This classification was a single-class labeling since each comment can only be *ham* or *spam* and has one polarity. After each coder classified all comments, we organized a virtual meeting to resolve all disagreements (863 regarding relevance and 1869 regarding polarity) with the two coders. In particular, we discussed each disagreement in the group to reach a consensus between the two coders.

Afterwards, we reviewed all 764 relevant (*ham*) comments to understand their intentions and topics [38]. We found the two main intentions *feature request* and *problem report* as well as the two main topics *efficiency* and *safety*. In the second classification, the same two coders classified all relevant comments regarding these two intentions and two topics. This classification was a multi-class labeling since these categories “are not mutually exclusive and can support each other when used in combination” [38, p. 8]. After each coder classified all comments, we organized again a virtual meeting to discuss and resolve all disagreements (191 regarding *feature request*, 378 regarding *problem report*, 246 regarding *efficiency*, and 158 regarding *safety*) in the group. We present the detailed results of the manual analysis in the sections IV-A and IV-B.

4) *Automated Analysis*: The automated analysis served to examine the *suitability* of video comments for current analysis techniques in CrowdRE. For this purpose, we decided to follow the overview of user feedback classification approaches [39].

³<https://developers.google.com/youtube/v3>

⁴<https://translate.google.com/>

Based on Santos et al. [39], we selected the three machine learning (ML) algorithms *support vector machines (SVM)*, *naive Bayes*, and *random forest*. While the first two algorithms are most frequently applied in CrowdRE, the third one is rarely used but showed very good performance [39]. We combined each of the three algorithms with the two ML features *bag of words (BOW)* and *term frequency - inverse document frequency (TF-IDF)*, which are most frequently used in CrowdRE [39]. The BOW feature included the use of only letters, tokenization, stemming, and the removal of stop words. For the implementation, we used *Python* and the library *scikit-learn*⁵ with the provided default settings of the particular algorithms. For each of these six combinations of algorithm and feature, we performed a 10-fold cross validation regarding the binary classification of *ham*, *feature request*, *problem report*, *efficiency*, and *safety*. Each 10-fold cross validation was repeated ten times to reduce the data splitting bias. We initially focused on binary classifiers since multi-class classifiers are even more complicated to use efficiently [37]. We present the detailed results of the automated analysis in section IV-C.

We published the raw and labeled data set as well as all additional materials, such as the guide sheets and *Python* scripts, online to increase the transparency of our results [40].

IV. RESULTS

Below, we present the results of analysis grouped by the three perspectives *quantity*, *content*, and *suitability*.

A. Quantity

This perspective focuses on how well the vision video can motivate stakeholders to participate in the crowd by watching the video and writing comments. As shown in TABLE I, the video had 6.9 million views and 4505 comments from 4302 different authors at the time we downloaded the data set (after almost 3.5 years). Over the next six months, the number of views, comments, and authors further increased. Although 4 years have passed after its release (see TABLE I), the video generated another 1.4 million views and 679 comments from 678 different authors. In CrowdRE, it is important to receive a large amount of feedback from many different stakeholders in a short period of time [17]. In Fig. 3, we illustrate the distribution of the number of comments and different authors over time. In the first four days after its release, the vision video solicited a total number of 2770 comments from 2660 different authors, which is 53.4% of all comments and 53.4% of all authors in four years. In the following 3 years, the number of comments increased slightly by 1142 (22.1%). In the fourth year, the video solicited 1272 comments (24.5%). Based on data set, we cannot explain this sharp increase in the last year.

Finding 1: The vision video motivated viewers to participate in the crowd. In a short period of time (four days), a large number of viewers watched the video and 2660 of them provided feedback by writing at least one comment.




Fig. 3: Number of comments and different authors over time

B. Content

This perspective considers the content of the comments and their relevance to RE. We summarized the results of the single- and multi-class labeling in TABLE II. Based on the single-class labeling, we identified 764 of the 4400 comments (17%) as relevant (*ham*) for RE. These comments include the authors’ opinions, questions, concerns, needs, and ideas combined with constructive feedback about the presented vision. The single-class labeling was restrictive since we determined that any comment that is promotional, incoherent, offensive, harmful, or lacked constructive feedback should be labeled as irrelevant (*spam*). As a result, the proportion of relevant comments is rather small. For example, the comment 3914 “*This Is Awesome.*” is a positive statement on the vision but lacks constructive feedback. Thus, the comment was classified as *spam*. In contrast, one author wrote the comment 517 “*Good idea but what about the possibility of sinkholes.*”. The comment was classified as *ham* since it provides constructive feedback by pointing out a potential problem.

TABLE II: Results of the labeling of the video comments

Category	Positive	Neutral	Negative	Overall
Overall	400	3033	967	4400
Spam	368	2486	782	3636
Ham	32	547	185	764
Feature request	7	115	20	142
Problem report	16	277	116	409
Efficiency	5	101	54	160
Safety	16	205	91	312

We further investigated the content of the relevant comments manually. In this way, we achieved a better understanding of the intentions of the comments. We found all three intentions *informing*, *reporting*, and *requesting* corresponding to the taxonomy for user feedback classification by Santos et al. [38]. *Informing* comments provide reasons based on features or problems either for or against the vision presented in the video, e.g., comment 151 “*A metro rail with extra steps, inefficiency, and fuel cost.*” *Reporting* comments address problems of the vision suspected by the author. One author wrote comment 271 “*Very unsafe I see no emergency exit at all.*” *Requesting*

⁵<https://scikit-learn.org/stable/>

comments focus on features of the vision. These comments either ask for missing functionality or share ideas for improving the vision by adding or changing features. For example, comment 302 “*it would be insane if you could charge your car while riding the tunnel*”. We simplified the multi-class labeling by initially focusing only on *feature request* and *problem report* as the two main intentions, similar to Maalej and Nabil [37].

In addition to the manual review, we used word clouds to explore frequent words and thus main topics in comments [28], [41]. In this way, we found that the two topics *efficiency* (see comment 151) and *safety* (see comment 271) are mainly addressed by the authors. These two intentions and these two topics form the categories for the multi-class labeling. For example, the following comment 1289 covers both intentions and both topics: “*As mentioned on other social media it may be a good idea to look into muti level above ground tunnels [(feature request)], although affected more by the weather eliments [(problem report)] would be easier to maintain [(efficiency)] and possibly much safer in the event of an earthquake or flooding [(safety)]*”.

Finding 2: The content of the comments is related to taxonomy of user feedback classification categories and thus relevant to RE. The relevant comments are mainly related to *feature requests* or *problem reports* regarding the presented vision. These comments primarily address the topics *efficiency* and *safety* of the envisioned system.

C. Suitability

This perspective investigates how well video comments are suitable for the current analysis techniques in CrowdRE. For this purpose, we performed an automated analysis of the video comments as described in section III-B4. In consideration of the highly varying numbers of comments in each category (see TABLE II), we decided to initially analyze a balanced data set for each category. Each balanced data set consists of as many data points as possible, e.g., the balanced data set on *safety* contains 312 comments on *safety* and 312 others randomly selected. TABLE III presents the average performance of the algorithms grouped by ML features and categories.

The analysis of all five categories shows promising results since the F_1 values range from 0.556 up to 0.821. However, the results obviously differ among the categories and algorithms.

1) *Categories:* While the relevant (*ham*) comments can be classified well by all algorithms ($0.657 < F_1 < 0.817$), the performance of the other classifications regarding the intentions and topics is divergent. The intention *problem report* consistently shows the lowest F_1 values (below 0.673) indicating difficulties in its classification. The other intention *feature request* and the topic *efficiency* have moderate results ($0.623 < F_1 < 0.745$). The topic *safety* is one of the best to classify. In case of using *SVM* with TF-IDF, we achieved the highest F_1 value of 0.821.

⁶The bold numbers highlight the best values per row. The conditional formatting is based on a comparison of all values per metric and the color gradient is from green (highest value) to yellow to red (lowest value).

TABLE III: Average performance of the algorithms based on a 10 times repeated 10-fold cross validation⁶

Category	Metric	Naive Bayes		SVM		Random forest	
		BOW	TF-IDF	BOW	TF-IDF	BOW	TF-IDF
Relevance: Ham	Accuracy	0.663	0.695	0.799	0.818	0.808	0.807
	Precision	0.675	0.697	0.799	0.819	0.809	0.808
	Recall	0.663	0.695	0.798	0.818	0.808	0.806
	F_1	0.657	0.694	0.798	0.817	0.807	0.806
Intention: Feature request	Accuracy	0.672	0.644	0.712	0.750	0.745	0.734
	Precision	0.680	0.656	0.726	0.757	0.753	0.742
	Recall	0.676	0.647	0.719	0.752	0.750	0.738
	F_1	0.666	0.634	0.708	0.745	0.742	0.731
Intention: Problem report	Accuracy	0.566	0.600	0.644	0.675	0.641	0.650
	Precision	0.570	0.600	0.651	0.681	0.644	0.656
	Recall	0.566	0.599	0.650	0.680	0.644	0.654
	F_1	0.556	0.594	0.604	0.673	0.640	0.648
Topic: Efficiency	Accuracy	0.642	0.632	0.723	0.737	0.747	0.724
	Precision	0.648	0.650	0.728	0.740	0.751	0.726
	Recall	0.643	0.637	0.726	0.740	0.746	0.725
	F_1	0.636	0.623	0.722	0.734	0.742	0.720
Topic: Safety	Accuracy	0.652	0.663	0.814	0.823	0.816	0.803
	Precision	0.670	0.667	0.822	0.827	0.822	0.806
	Recall	0.651	0.664	0.815	0.824	0.816	0.805
	F_1	0.638	0.658	0.811	0.821	0.813	0.801

2) *Algorithms:* *Naive Bayes* consistently shows the lowest results for all categories ($0.556 < F_1 < 0.666$). The results of *SVM* ($0.604 < F_1 < 0.821$) and *random forest* ($0.640 < F_1 < 0.813$) are always better than those of *naive Bayes* for each category. *SVM* and *random forest*, in turn, achieve very similar results. Nevertheless, in almost all cases, *SVM* with TF-IDF achieves slightly better results than *random forest* (see TABLE III, bold highlighting). Only for the topic *efficiency* are the results of *random forest* with BOW slightly better.

These observations indicate that video comments are suitable for current analysis techniques in CrowdRE. Already the use of algorithms with their default settings (see section III-B4) leads to first promising results. Even though the algorithms produce different results, the use of *SVM* with TF-IDF shows the best results overall.

Finding 3: The video comments are suitable for current analysis techniques in CrowdRE. Despite differences in the performance of the algorithms used and the individual categories, the automated analysis shows promising results. While *problem reports* are more difficult to classify, especially relevant (*ham*) comments and comments on *safety* can be classified well. In particular, the use of *SVM* with TF-IDF achieved the best results overall.

V. THREATS TO VALIDITY

Below, we report threats to construct, external, internal, and conclusion validity according to Wohlin et al. [42].

A. Construct Validity

There is a mono-operation bias since the study is based on one case. Although the case study is a real-world example, one case cannot represent the complexity in reality. In addition, we must emphasize that this case is special since the vision

video comes from a company founded by Elon Musk and already has a strong social media presence. The selection of three ML algorithms and two features leads to a method bias (see section III-B4). According to Santos et al. [39], there are at least six other algorithms and five other features that are used in CrowdRE. We assume that our selection is a good compromise for an initial study since we focus on the two most frequently used algorithms and features. In addition, we use an algorithm that is rarely used but has shown a very good performance [39]. However, the selection is necessary since this study first investigates whether the use of vision videos on social media platforms has any potential to solicit feedback in the form of video comments for CrowdRE. For this reason, we also neglected the tuning of the hyperparameters and used the algorithms with their default settings. As a result, this case study is only a preliminary investigation and future replications and extensions are needed to strengthen the results and draw more reliable conclusions.

B. External Validity

The single case limits the generalization of the results. However, generalization is not the focus of this study. Instead, we want to get a first impression of the potential of using vision videos on social media platforms for CrowdRE. First, we attempted to find existing labeled data sets that we could have used to increase external validity. However, all of the data sets we found only contained labeled video comments on coding tutorials, music videos, or movie trailers and thus did not fit our context and purpose. We also contacted Vistisen and Poulsen [5] to obtain their data set. Unfortunately, they only created a data set that contains the raw data without the labels. For future extensions of our case study, we currently create additional data sets based on the raw data set of Vistisen and Poulsen [5] as well as another vision video⁷ from YouTube.

C. Internal Validity

The video comments often have a poor text quality and belong to several categories, even beyond those categories classified in this study, such as *user experience* and *ratings* found by Maalej and Nabil [37]. These shortcomings impede the manual single- and multi-class labeling. We provided the coders with a guide sheet for their orientation to mitigate this threat to validity [43]. After each labeling, we calculated *Cohen's κ* [44] for each category (see TABLE IV). We found a maximum value of 0.35, indicating only moderate agreement [45] and thus high subjectivity of the classification by the two coders. For this reason, we organized two meetings to discuss each disagreement in the group and reach a consensus between the two coders. As suggested by Maalej and Nabil [37], this procedure helps to resolve disagreements and reach a consensus. Nevertheless, we cannot guarantee that all video comments are completely and correctly classified. Furthermore, we simplified the labeling by focusing only on five categories. According to the taxonomy of user feedback classification

categories [38], there are many more (sub-) categories. However, we neglected these categories for now due to the preliminary nature of our study.

TABLE IV: Inter-rater reliability between the two coders

Category	Relevance	Polarity	Feature request	Problem report	Efficiency	Safety
Cohen's κ	0.35	0.14	0.15	-0.01	0.28	0.26

D. Conclusion Validity

All our findings are based on the results of the manual analysis and thus the subjective interpretation of the video comments. We cannot completely exclude the misinterpretation of some video comments due to a missing context. Nevertheless, we mitigated this threat to validity with a defined strategy. While the first two authors planned the data analysis procedure, the third author reviewed the design and suggested improvements. Afterwards, the second author performed the entire data analysis procedure with two external coders, the first author supervised the procedure, and the third author reviewed the analysis and results [46]. For this reason, we are confident that the findings are valid and sound.

VI. DISCUSSION

The case study provides important insights into the potential of using vision videos on social media platforms to solicit feedback in form of video comments for CrowdRE. Our findings substantiate that vision videos (1) can solicit many comments from various viewers in a short period of time, which (2) contain relevant content to RE, and (3) are suitable for current analysis techniques in CrowdRE.

The use of vision videos on social media platforms shows promise for soliciting feedback in form of video comments. We investigated the potential of using vision videos for CrowdRE from the three perspectives *quantity*, *content*, and *suitability*. The manual analysis shows that the use of the vision video on YouTube resulted in thousands of views and comments within a few days (see Finding 1). Although we found that only a certain fraction of these comments are relevant to RE (see Finding 2), these comments address typical intentions and topics of user feedback, such as *feature request*, *problem report*, and *efficiency*. Despite the lower fraction of relevant comments, automated analyses are unavoidable due to the large number of comments. The automated analysis shows that video comments are suitable to be analyzed by ML algorithms (frequently) used in CrowdRE. Already algorithms with their default settings achieved a good performance (see Finding 3). For these reasons, we conclude that using vision videos on social media platforms offers great potential to solicit video comments as a source of feedback for CrowdRE.

Most of our findings correspond with the current state of knowledge in CrowdRE literature. However, we also found one new insight. While the topic *safety* is addressed in over 300 video comments, the taxonomy of user feedback classification categories does not contain this characteristic in the *quality*

⁷<https://www.youtube.com/watch?v=S5fOWB6SNqs>

in use category [38]. Our topic *safety* corresponds to the characteristic *freedom of risk* of ISO/IEC FDIS 25010:2010 [47] addressing economic, health, safety, and environmental risks. We assume that this difference is a result of using a vision video as stimulus showing a cyber-physical system. Previous studies have primarily focused on user feedback on pure software systems leading to a taxonomy for user feedback classification [38] that does not include *safety* as a characteristic. Vision videos can stimulate this kind of feedback since they are concrete by showing specific scenarios and contexts [14]. Thus, vision videos do not only show the software system but also its use in concrete situations. Viewers can more easily put themselves in these concrete situations [4], [48], [49]. In this way, they experience the system before it has even been developed. Through the use of vision videos, we are able to empower stakeholders to report their suspected *quality in use* problems, such as economic, health, safety, and environmental risks, even if they have never actually used the system.

The vision video used in this case study is one important point we need to discuss. The video comes from a company founded by Elon Musk and already has a strong social media presence. As a result, this case is special since the conditions to receive feedback were much better than for average companies. Successfully soliciting feedback in the form of video comments requires a good strategy for building a corresponding social media presence for a company. This strategy must go beyond marketing for bridging the gap to the development so that a cycle can be created in the sense of CrowdRE to solicit, analyze, implement, and validate feedback in order to solicit new feedback again. For this purpose, future research must develop systematic and holistic approaches aligned with such a strategy and covering the entire software development process to operationally involve a crowd, e.g., to provide content, form a social network, and support the development with feedback.

For the analysis of the feedback, we currently used traditional and established machine learning algorithms and features. Even though the results are positive, there is a new trend in text classification using embeddings and transformer-based classifiers such as BERT [51]. First approaches using deep learning algorithms such as BERT show promising results in classifying user feedback [52]. Our results show that video comments are another source of feedback whose analysis needs to be further developed, especially considering this current research trend. As an answer to our research question, we can summarize:

Answer: The use of vision videos on social media platforms offers a large potential to solicit feedback in form of video comments for CrowdRE. According to the manual and automated analysis, a sufficient number of comments emerge whose content is relevant to RE and which are suitable for current analysis techniques in CrowdRE. The findings point to the conclusion that using vision videos for CrowdRE can motivate stakeholders to actively participate in a crowd by writing video comments that are a valuable source of feedback.

VII. CONCLUSION & FUTURE WORK

In this paper, we present the findings of an initial case study on the potential of using of vision videos for CrowdRE. According to these findings, the use of vision videos for CrowdRE has a large potential since these videos can solicit a large amount of feedback from various stakeholders in a short period of time. In addition to typical user feedback categories, such as *feature request* and *problem report*, we found that vision videos empower viewers to experience the future system in a way that allows them to comment on suspected problems related to *safety*. This finding is particularly interesting since this topic is an entirely new category that has not previously appeared in analyses of user feedback. However, the findings should not be overgeneralized due to the preliminary nature of our case study. Nevertheless, we are optimistic that vision videos can motivate stakeholder to actively participate in a crowd, solicit feedback, and stimulate discussion for CrowdRE. While our case study substantiates the first two points, the stimulation of discussions is only an assumption based on a manual review of the replies. Similar to Vistisen and Poulsen [5], we have the impression that replies do not address the video itself but primarily discuss the associated comment. This assumption needs to be investigated in the future, for example with the *CrowdRE-Arg* framework by Khan et al. [53]. This framework can identify arguments that are for or against a given statement. Overall, the findings show a positive trend, but future replications and extensions of this study are needed for a more comprehensive assessment of the potential of using vision videos for CrowdRE.

REFERENCES

- [1] A. Bennaceur, C. McCormick, J. García-Galán, C. Perera, A. Smith, A. Zisman, and B. Nuseibeh, "Feed me, Feed me: An Exemplar for Engineering Adaptive Software," in *2016 IEEE/ACM 11th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*. IEEE, 2016, pp. 89–95.
- [2] A. Darby, E. Tsekleves, and P. Sawyer, "Speculative Requirements: Design Fiction and RE," in *2018 IEEE 26th International Requirements Engineering Conference (RE)*. IEEE, 2018, pp. 388–393.
- [3] K. Schneider, M. Busch, O. Karras, M. Schrapel, and M. Rohs, "Refining Vision Videos," in *International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ)*. Springer, 2019.
- [4] M. Busch, O. Karras, K. Schneider, and M. Ahrens, "Vision Meets Visualization: Are Animated Videos an Alternative?" in *International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ)*. Springer, 2020, pp. 277–292.
- [5] P. Vistisen and S. B. Poulsen, "Return of the Vision Video: Can Corporate Vision Videos Serve as Setting for Participation?" in *7th Nordic Design Research Conference*, no. 7, 2017.
- [6] O. Karras, K. Schneider, and S. A. Fricker, "Representing Software Project Vision by Means of Video: A Quality Model for Vision Videos," *Journal of Systems and Software*, vol. 162, 2020.
- [7] O. Karras, *Supporting Requirements Communication for Shared Understanding by Applying Vision Videos in Requirements Engineering*. Logos Verlag Berlin GmbH, 2021.
- [8] K. Schneider, O. Karras, A. Finger, and B. Zibell, "Reframing Societal Discourse as Requirements Negotiation: Vision Statement," in *2017 IEEE 25th International Requirements Engineering Conference Workshops (REW)*. IEEE, 2017, pp. 188–193.
- [9] S. W. Ambler, *Agile Modeling: Effective Practices for eXtreme Programming and the Unified Process*. John Wiley & Sons, 2002.
- [10] O. Karras and K. Schneider, "Software Professionals are Not Directors: What Constitutes a Good Video?" in *2018 1st International Workshop on Learning from other Disciplines for Requirements Engineering (DARE)*. IEEE, 2018, pp. 18–21.

- [11] O. Karras, A. Hamadeh, and K. Schneider, "Enriching Requirements Specifications with Videos – The Use of Videos to Support Requirements Communication," *GI Softwaretechnik-Trends*, vol. 38, no. 1, pp. 51–52, 2017.
- [12] O. Karras, C. Unger-Windeler, L. Glauer, and K. Schneider, "Video as a By-Product of Digital Prototyping: Capturing the Dynamic Aspect of Interaction," in *2017 IEEE 25th International Requirements Engineering Conference Workshops (REW)*. IEEE, 2017, pp. 118–124.
- [13] O. Karras, "Software Professionals' Attitudes towards Video as a Medium in Requirements Engineering," in *International Conference on Product-Focused Software Process Improvement*. Springer, 2018, pp. 150–158.
- [14] O. Karras, J. Klünder, and K. Schneider, "Enrichment of Requirements Specifications with Videos: Enhancing the Comprehensibility of Textual Requirements," in *Proceedings of Videos in Digital Libraries – What's in it for Libraries, Scientists, and Publishers? (TPDL)*. Zenodo, 2016.
- [15] K. Schneider and L. M. Bertolli, "Video Variants for CrowdRE: How to Create Linear Videos, Vision Videos, and Interactive Videos," in *2019 IEEE 27th International Requirements Engineering Conference Workshops (REW)*. IEEE, 2019, pp. 186–192.
- [16] J. A. Khan, L. Liu, L. Wen, and R. Ali, "Crowd Intelligence in Requirements Engineering: Current Status and Future Directions," in *International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ)*. Springer, 2019, pp. 245–261.
- [17] E. C. Groen and M. Koch, "How Requirements Engineering can Benefit from Crowds," *Requirements Engineering Magazine*, vol. 8, p. 10, 2016.
- [18] M. Stade, M. Oriol, O. Cabrera, F. Fotrousi, R. Schaniel, N. Seyff, and O. Schmidt, "Providing a User Forum is not Enough: First Experiences of a Software Company with CrowdRE," in *2017 IEEE 25th International Requirements Engineering Conference Workshops (REW)*. IEEE, 2017.
- [19] C. Wang, M. Daneva, M. van Sinderen, and P. Liang, "A Systematic Mapping Study on Crowdsourced Requirements Engineering Using User Feedback," *Journal of Software: Evolution and Process*, vol. 31, no. 10, 2019.
- [20] S. Lim, A. Henriksson, and J. Zdravkovic, "Data-Driven Requirements Elicitation: A Systematic Literature Review," *SN Computer Science*, vol. 2, no. 1, pp. 1–35, 2021.
- [21] O. Karras, S. Polst, and K. Späth, "Using Vision Videos in a Virtual Focus Group: Experiences and Recommendations," *GI Softwaretechnik-Trends*, vol. 41, no. 1, 2021.
- [22] G. Broll, H. Hussmann, E. Rukzio, and R. Wimmer, "Using Video Clips to Support Requirements Elicitation in Focus Groups – An Experience Report," in *SE Workshop on Multimedia Requirements Engineering*, 2007.
- [23] O. Karras, S. Kiesling, and K. Schneider, "Supporting Requirements Elicitation by Tool-Supported Video Analysis," in *2016 IEEE 24th International Requirements Engineering Conference (RE)*. IEEE, 2016, pp. 146–155.
- [24] E. C. Groen, J. Schowalter, S. Kopczynska, S. Polst, and S. Alvani, "Is there Really a Need for Using NLP to Elicit Requirements? A Benchmarking Study to Assess Scalability of Manual Analysis," in *REFSQ Workshops*, 2018.
- [25] E. C. Groen, "Crowd out the Competition," in *2015 IEEE 1st International Workshop on Crowd-Based Requirements Engineering (CrowdRE)*. IEEE, 2015, pp. 13–18.
- [26] S. Sharmin and Z. Zaman, "Spam Detection in Social Media Employing Machine Learning Tool for Text Mining," in *13th International Conference on Signal-Image Technology & Internet-Based Systems*. IEEE, 2017, pp. 137–142.
- [27] A. O. Abdullah, M. A. Ali, M. Karabatak, and A. Sengur, "A Comparative Analysis of Common YouTube Comment Spam Filtering Techniques," in *6th International Symposium on Digital Forensic and Security*. IEEE, 2018, pp. 1–5.
- [28] A. Obadimu, E. Mead, M. N. Hussain, and N. Agarwal, "Identifying Toxicity within YouTube Video Comment," in *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*. Springer, 2019.
- [29] R. K. Das, S. S. Dash, K. Das, and M. Panda, "Detection of Spam in YouTube Comments Using Different Classifiers," in *Advanced Computing and Intelligent Engineering*. Springer, 2020, pp. 201–214.
- [30] M. Z. Asghar, S. Ahmad, A. Marwat, and F. M. Kundi, "Sentiment Analysis on YouTube: A Brief Survey," *arXiv preprint arXiv:1511.09142*, 2015.
- [31] A. U. R. Khan, M. Khan, and M. B. Khan, "Naïve Multi-label Classification of YouTube Comments Using Comparative Opinion Mining," *Procedia Computer Science*, vol. 82, pp. 57–64, 2016.
- [32] E. Poché, N. Jha, G. Williams, J. Staten, M. Vesper, and A. Mahmoud, "Analyzing User Comments on YouTube Coding Tutorial Videos," in *IEEE/ACM 25th International Conference on Program Comprehension*. IEEE, 2017, pp. 196–206.
- [33] R. Benkhelifa and F. Z. Laallam, "Opinion Extraction and Classification of Real-Time YouTube cooking Recipes Comments," in *International Conference on Advanced Machine Learning Technologies and Applications*. Springer, 2018, pp. 395–404.
- [34] P. Runeson and M. Höst, "Guidelines for Conducting and Reporting Case Study Research in Software Engineering," *Empirical Software Engineering*, vol. 14, no. 2, p. 131, 2009.
- [35] J. Burgess and J. Green, *YouTube: Online Video and Participatory Culture*. John Wiley & Sons, 2018.
- [36] C. Schwemmer and S. Ziewiecki, "Social Media Sellout: The Increasing Role of Product Promotion on YouTube," *Social Media+ Society*, vol. 4, no. 3, 2018.
- [37] W. Maalej and H. Nabil, "Bug Report, Feature Request, or simply Praise? On Automatically Classifying App Reviews," in *2015 IEEE 23rd International Requirements Engineering Conference (RE)*. IEEE, 2015.
- [38] R. Santos, E. C. Groen, and K. Villela, "A Taxonomy for User Feedback Classifications," in *REFSQ Workshops*, 2019.
- [39] —, "An Overview of User Feedback Classification Approaches," in *REFSQ Workshops*, 2019.
- [40] O. Karras and E. Kristo, "Dataset of Video Comments of a Vision Video Classified by Their Relevance, Polarity, Intention, and Topic." Zenodo, 2021. [Online]. Available: <https://doi.org/10.5281/zenodo.4698969>
- [41] F. Heimerl, S. Lohmann, S. Lange, and T. Ertl, "Word Cloud Explorer: Text Analytics Based on Word Clouds," in *47th Hawaii International Conference on System Sciences*. IEEE, 2014, pp. 1833–1842.
- [42] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in Software Engineering*. Springer Science & Business Media, 2012.
- [43] K. A. Neuendorf, *The Content Analysis Guidebook*. Sage Publications, 2015.
- [44] J. Cohen, "A Coefficient of Agreement for Nominal Scales," *Educational and Psychological Measurement*, vol. 20, no. 1, pp. 37–46, 1960.
- [45] J. R. Landis and G. G. Koch, "The Measurement of Observer Agreement for Categorical Data," *Biometrics*, pp. 159–174, 1977.
- [46] E. Kristo, "Computer-Aided Analysis of Video Comments for Requirements Analysis," Bachelor thesis, Leibniz Universität Hannover, 2021.
- [47] "ISO/IEC FDIS 25010:2010: Systems and Software Engineering – Systems and Software Product Quality Requirements and Evaluation (SQuaRE) – System and Software Quality Models," 2010.
- [48] A. Ferrari, P. Spoletini, and S. Gnesi, "Ambiguity as a Resource to Disclose Tacit Knowledge," in *2015 IEEE 23rd International Requirements Engineering Conference (RE)*. IEEE, 2015, pp. 26–35.
- [49] O. Liskin, "How Artifacts Support and Impede Requirements Communication," in *International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ)*. Springer, 2015, pp. 132–147.
- [50] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks," *arXiv preprint arXiv:1908.10084*, 2019.
- [51] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-Training of Deep Bidirectional Transformers for Language Understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [52] R. Reddy Mekala, A. Irfan, A. Porter, M. Lindvall, and E. C. Groen, "Classifying User Requirements from Online Feedback in Small Dataset Environments Using Deep Learning," in *2021 IEEE 29th International Requirements Engineering Conference*. IEEE, 2021.
- [53] J. Ali Khan, L. Liu, L. Wen, and R. Ali, "Conceptualising, Extracting and Analysing Requirements Arguments in Users' Forums: The CrowdRE-Arg Framework," *Journal of Software: Evolution and Process*, vol. 32, no. 12, 2020.