



Validating Multimedia Content Moderation Software via Semantic Fusion

Wenxuan Wang
The Chinese University of Hong Kong
Hong Kong, China
wxwang@cse.cuhk.edu.hk

Jingyuan Huang
The Chinese University of Hong Kong
Hong Kong, China
1155173905@link.cuhk.edu.hk

Chang Chen
The Chinese University of Hong Kong
Hong Kong, China
ccchen@link.cuhk.edu.hk

Jiazhen Gu
The Chinese University of Hong Kong
Hong Kong, China
jiazheng@cuhk.edu.hk

Jianping Zhang
The Chinese University of Hong Kong
Hong Kong, China
jpzhang@cse.cuhk.edu.hk

Weibin Wu
Sun Yat-sen University
Zhuhai, China
wuwb36@mail.sysu.edu.cn

Pinjia He*
The Chinese University of Hong
Kong, Shenzhen
Shenzhen, China
hepinjia@cuhk.edu.cn

Michael Lyu
The Chinese University of Hong Kong
Hong Kong, China
lyu@cse.cuhk.edu.hk

ABSTRACT

The exponential growth of social media platforms, such as Facebook, Instagram, Youtube, and TikTok, has revolutionized communication and content publication in human society. Users on these platforms can publish multimedia content that delivers information via the combination of text, audio, images, and video. Meanwhile, the multimedia content release facility has been increasingly exploited to propagate toxic content, such as hate speech, malicious advertisement, and pornography. To this end, content moderation software has been widely deployed on these platforms to detect and blocks toxic content. However, due to the complexity of content moderation models and the difficulty of understanding information across multiple modalities, existing content moderation software can fail to detect toxic content, which often leads to extremely negative impacts (e.g., harmful effects on teen mental health).

We introduce *Semantic Fusion*, a general, effective methodology for validating multimedia content moderation software. Our key idea is to fuse two or more existing single-modal inputs (e.g., a textual sentence and an image) into a new input that combines the semantics of its ancestors in a novel manner and has toxic nature by construction. This fused input is then used for validating multimedia content moderation software. We realized *Semantic Fusion* as *DUO*, a practical content moderation software testing tool. In our evaluation, we employ *DUO* to test five commercial content moderation software and two state-of-the-art models against three

kinds of toxic contents. The results show that *DUO* achieves up to 100% error finding rate (EFR) when testing moderation software and it obtains up to 94.1% EFR when testing the state-of-the-art models. In addition, we leverage the test cases generated by *DUO* to retrain the two models we explored, which largely improves model robustness (2.5%~5.7% EFR) while maintaining the accuracy on the original test set.

CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging**.

KEYWORDS

Software testing, semantic fusion, multimedia content moderation, metamorphic testing

ACM Reference Format:

Wenxuan Wang, Jingyuan Huang, Chang Chen, Jiazhen Gu, Jianping Zhang, Weibin Wu, Pinjia He, and Michael Lyu. 2023. Validating Multimedia Content Moderation Software via Semantic Fusion. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA '23)*, July 17–21, 2023, Seattle, WA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3597926.3598079>

1 INTRODUCTION

Multimedia contents, such as Internet memes and videos, play an important role in online communication and content publication on social media platforms. For example, in 2020, there were more than one million posts mentioning "meme" being shared on Instagram, one of the most popular social media platforms [3], every day. Moreover, Cisco reports that 82% of global Internet traffic will come from either video streaming or video downloads in 2022 [55]. Fig. 1 presents two examples of such multimedia content from the web^{1,2}. Although the exponential growth of multimedia content has

*Pinjia He is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ISSTA '23, July 17–21, 2023, Seattle, WA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0221-1/23/07...\$15.00
<https://doi.org/10.1145/3597926.3598079>

¹<https://www.pinterest.com/pin/top-phd-memes-of-2020-thephdhub-400327854385638920/>

²<https://www.pinterest.com/pin/tv-3-359302876503124197/>



Figure 1: Examples of multimedia content from the web: (1) a meme and (2) a video frame with subtitles.

greatly facilitated user communication and content distribution in the world, it has also exacerbated the propagation of *toxic content*.

Toxic contents generally refer to harmful contents that can cause negative affect on reader's attitudes, behavior or health. In particular, toxic contents are mainly including but not limited to into the following categories: (1) *abusive language and hate speech*, which are abusive contents targeting specific individuals, such as politicians, celebrities, religions, nations, and the LGBTIQ+ [1]; (2) *malicious advertisements*, which are online advertisements with illegal purposes, including phishing and scam links, malware download, and illegal information dissemination [36]; and (3) *pornography*, which is often sexually explicit, associative, and aroused [49].

Such toxic contents have significant negative impacts on users. For instance, Munro [43] concluded that online *hate speech* may develop depression, anxiety, and other mental health problems in children. *Malicious advertisements* for illegal purposes also remain a global burden, accounting for up to 85% of daily message traffic [8]. *Pornography* may cause significant undesirable effects on both the physical and psychological health of children [76]. Statistics showed that adult content sites accounted for 0.67% of all website categories accessed by Latin American children from May 2019 to May 2020 [9]. Moreover, such widely disseminated toxic contents greatly affect social harmony and increase the number of criminal cases to a certain extent [6].

Due to the harmfulness of toxic content, content moderation software for detecting and blocking toxic content has attracted massive interest from both academia and industry. Existing methods typically formulate toxic content detection as a classification task and resort to Artificial Intelligence (AI) techniques, such as convolutional neural networks[31], long-short-term-memory models[20], and Transformer-based models[65], and achieving considerable performance on corresponding datasets[41, 53, 73]. Because of the importance of content moderation, large-scale online service providers, such as Google [17], Meta (Facebook) [66], Twitter [14], and Baidu [26], have extensively deployed commercial-level content moderation software on their products. In particular, Meta reports that they remove millions of violating content on Facebook

and Instagram every day, among which more than 90% are detected by AI-based moderation software³.

Although tremendous efforts have been spent on developing toxic content moderation models, existing content moderation software sometimes fails to detect inputs from malicious users, exposing toxic content to other users. For example, a New Zealand terrorist live-streamed a massacre on Facebook [57]. YouTube Kids, an app for children, was reported to contain a significant amount of inappropriate content, which was made readily available for unsuspecting kids [33]. Due to the huge number of Internet users, even a 0.01% failure rate may cause serious consequences.

Despite its apparent importance, validating the robustness of multimedia content moderation software is very difficult and has, therefore, been much under-explored. First, existing high-quality multimedia data have already been utilized in the development of the moderation software and models, while the construction of a new test oracle typically incurs extensive manual labeling effort. Second, previous studies mainly generate test cases from only one specific modality, such as visual modality [77], audio modality [22], and textual modality [47, 48]. While these approaches can generate interesting test cases, they fail to stress-test the core ability of multimedia content moderation software: understanding multi-modal inputs.

Inspired by an SMT solver validation method [71] that fuses two existing formulas into a new formula, this paper introduces *Semantic Fusion*, a general, effective methodology for validating multimedia content moderation software. Our key insight is to fuse two or more single-modal inputs into a new multi-modal input that combines the semantics of its ancestors and is toxic by construction. Fig. 2 presents the high-level idea of *Semantic Fusion* via two sketches of the test cases. To realize this concept, we implement *DUO*, a tool that can generate test cases that cover all three typical categories of toxic contents (*i.e.*, hate speech, malicious advertisement, and pornography) in two widely-used languages (*i.e.*, English and Chinese). Specifically, *DUO* first adopts a template-based approach to construct seed toxic sentences. Then *DUO* generates multi-modal toxic contents as test cases by distributing the information of toxic sentences into different modalities and fusing single-modal inputs accordingly. These toxic contents will be fed into the multimedia content moderation software as test cases. If a test case evades the detection of the software under test, an error will be reported.

To evaluate the performance of *DUO*, we apply *DUO* to test five widely-deployed commercial content moderation software from famous software providers, including Google Cloud, Amazon Web Service, Baidu Cloud, Tencent Cloud, and Alibaba Cloud, and two state-of-the-art models (*i.e.*, Vision-Transformer-based [12] and ResNet-based models [18]). The results show that the software under test fails to detect most of the test cases generated by *DUO*. Notably, up to 100% and 94.1% of generated toxic image and video test cases can bypass the content moderation software and research models, respectively. In addition, we leverage the test cases generated by *DUO* to retrain the model we explored, which largely improves

³<https://transparency.fb.com/zh-cn/enforcement/detecting-violations/technology-detects-violations/>

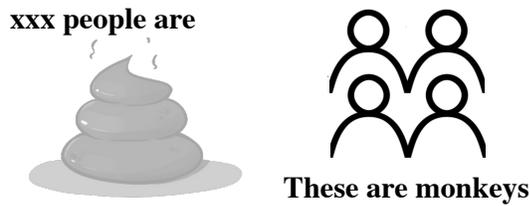


Figure 2: The sketches of two multimedia contents generated by *Semantic Fusion*. Both contents are toxic (hate speech) by construction if: (1) "xxx" was replaced by the name of a specific group (e.g., racial, gender, religion, or nation) for the left content; and (2) the people were replaced by a photo of a specific group for the right content.

model robustness (2.5%~5.7% EFR) while maintaining the accuracy on the original test set.

The contributions of this paper are summarized as follows:

- We introduce *Semantic Fusion*, a general, effective methodology for testing multimedia content moderation software.
- Based on the *Semantic Fusion* methodology, we design and develop the first tool, DUO, for multimedia content moderation software validation.
- DUO effectively reported errors in five widely-deployed commercial software products and two state-of-the-art research models with consistently high error finding rates.
- We successfully improved the robustness of the two moderation models by retraining them with the failed test cases.

Content Warning: We apologize that this article presents examples of aggressive, abusive, and pornographic expressions to demonstrate the results of our method.

2 BACKGROUND

2.1 Content Moderation Software

Big companies, for example, Google, Meta, Twitter, Amazon, Baidu, Tencent and Alibaba, have developed and deployed commercial-level content moderation software on their products. According to their official technical documents, the backbone of their software is usually a completed engineering system containing neural network-based models and rule-based methods. For example, Baidu Commercial Content Moderation Software is powered by a deep neural network and a huge pre-defined banned word list. This kind of hybrid approach can leverage the best of different methods. Neural network-based methods can effectively mine contexts and semantic information, while rule-based methods easily implement user-defined functionality.

2.2 Multi-modal Multimedia Content

Multimedia is a form of communication that uses a combination of different content forms such as text, audio, images, animations, or video into a single presentation⁴. A model that can deal with multimedia data is called multi-modal AI and each channel is called

a modality. There are two main issues in multi-modal AI, processing single modality and understanding information across different modalities. Our work mainly involves three modalities: visual modality, audio modality, and textual modality. For example, a meme contains visual (the image) and textual (the top text and bottom text) information, and a video usually involves visual modality (the video screen), audio modality (the soundtrack), and textual modality (the subtitles).

2.3 Multi-modal Fusion

Multimedia content (e.g., a meme or video) has different modalities to convey information. Therefore, to understand the whole picture of multimedia content and determine its toxicity, one needs not only to process the information in every single modality but also to fuse the information from different modalities. The fusion of different modalities is generally performed at two levels: feature level and decision level. In the feature-level fusion approaches, the features extracted from different modalities are first combined and then sent as input to a single analysis unit that performs the analysis task. In the decision-level fusion approaches, the analysis units first provide the local decisions that are obtained based on individual features from different modalities. The local decisions are then combined using a decision fusion unit to make a fused decision. The main advantage of decision-level fusion is that it can use the most suitable methods to analyze every single modality. However, it fails to utilize the feature-level correlation among modalities.

3 APPROACH AND IMPLEMENTATION

In this section, we introduce the design and implementation of DUO, a novel tool to validate content moderation software. Figure 3 overviews the workflow of DUO, which consists of two main modules, test case generation and error detection. In particular, the test case generation module adopts a semantic fusion-based method to generate toxic multimedia contents as test cases. Then, the generated cases are fed into the error detection module, which performs metamorphic testing to reveal errors in content moderation software.

More specifically, to generate toxic multimedia content, DUO first adopts a template-based method to extract keyword pairs from existing toxic datasets. Each keyword pair can constitute toxic sentences, which are further used as the seeds to generate toxic multimedia contents (Section 3.1). For each extracted keyword pair, DUO distributes the keyword pairs into different modalities (Section 3.2), and then fuses the multi-modal information to generate multimedia contents (Section 3.3). Finally, DUO feeds the generated cases into the content moderation software under test, and detects errors based on metamorphic relations, i.e., the content toxicity is invariant under modality transformation (Section 3.4).

3.1 Keyword Pair Extraction

To generate test cases, DUO first extracts keyword pairs, which can constitute toxic contents, from existing datasets. Specifically, DUO adopts a template-based approach: It first extracts keyword pairs that can generate toxic sentences for different toxic categories based on templates. Then, an NLP-based filtering method is adopted

⁴<https://en.wikipedia.org/wiki/Multimedia>

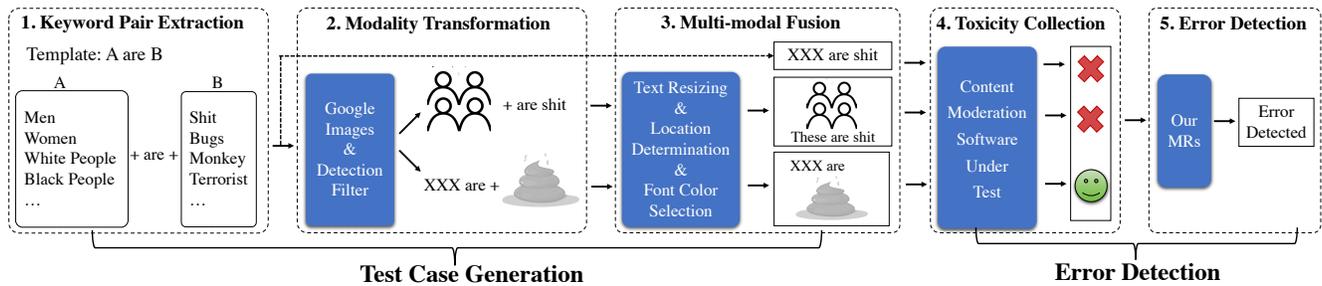


Figure 3: The overview of DUO.

to drop invalid keyword pairs, so as to avoid the generation of low-quality cases.

3.1.1 Template Designing. DUO utilizes the following templates to extract keyword pairs for each kind of toxicity, *i.e.*, hate, advertisement and pornography.

For abuse and hate speech, DUO utilizes the template "A is/are B", where A is a group, *e.g.*, a specific race or gender, and B is a negative adjective, such as "stupid" or "lazy", or noun, such as criminal or pig.

For malicious advertisement, we design a template of "A: B", where A is a product, such as "Tobacco" or "Alcohol" and B is the contact information, such as telephone number, email address or WhatsApp number.

For pornography, the template "A your/my B" is adopted, where A is a verb and B is a sexual-related organ.

Even though DUO only adopts one simple template for each type of toxicity, it can detect a large number of errors in practical content moderation software (Section 4). Since DUO can achieve satisfying performance even with simple templates, we do not investigate the effectiveness of more complicated templates, which can be studied in future work.

3.1.2 Keyword Selection. Based on the above templates for different kinds of toxicity, DUO extracts keyword pairs that can constitute corresponding toxic texts from existing toxic datasets. In order to find keywords similar to real-world cases, DUO utilizes a total of 6 manually labeled datasets collected from practical Internet platforms, 2 for each toxicity type. The statistics of the datasets are shown in Table 1.

Specifically, for abuse and hate speech, DUO extracts keyword pairs from Social Bias Corpus [52] and Dirty⁵. Social Bias Corpus contains 150k structured annotations of social media posts, covering over 34k implications about a thousand demographic groups. Dirty is an open GitHub repository containing 2.5k Chinese toxic sentences with abusive and sexual words.

For malicious advertisement, DUO extracts keyword pairs from SMS Spam Collection⁶ and SpamMessage⁷. SMS Spam Collection is a set of tagged SMS messages, containing 5,574 SMS messages in English, tagged as being ham (legitimate) or spam. The data was

Table 1: Statistics of toxic datasets.

Dataset	#Sent	Lang	Type	Source
Social Bias	150k	English	Abuse	Twitter
Dirty	2.5K	Chinese	Abuse	Weibo
SMSSpam	5.5k	English	Ad.	Grumbletext
SpamMessage	60K	Chinese	Ad.	Taobao
Sexting	0.5K	English	Porno	Github
Midu	7.3K	Chinese	Porno	Midu

manually extracted from the Grumbletext website, a UK forum in which cell phone users make public claims about SMS spam messages, while SpamMessage is an open GitHub repository containing 60k malicious advertisement messages.

For pornography, DUO utilizes Sexting⁸ and Midu [58], where Sexting is an English pornographic text dataset containing 537 sexual texting messages, while Midu is a Chinese novel paragraph dataset collected by ourselves from an online literature reading platform called MiDu App⁹. It is a corpus with 62,876 paragraphs including 7,360 pornographic paragraphs and 55,516 normal paragraphs.

It is worth noting that not all words in the above datasets are potential keywords. An ideal keyword should be frequently used in toxic content while less frequently in a general domain corpus so that it is more likely to contain toxicity. Therefore, we use TF-IDF, a numerical statistic that reflects how important a word is to a document in a collection or corpus, to select potential keywords from the above datasets. The TF-IDF value increases proportionally to the number of times a word appears in the document and is offset by the number of documents in the corpus that contain the word. Particularly, DUO utilizes sklearn¹⁰ (English) and Jieba library¹¹ (Chinese) to filter out the stop words, followed by calculating the TF-IDF score, and select the top 20 words with the highest TF-IDF score as the candidate keywords for each dataset.

3.1.3 Keyword Pair Extraction and Filtering. After obtaining the candidate keywords, it is non-trivial to extract keyword pairs that

⁵<https://github.com/pokemonchw/Dirty>

⁶<https://www.kaggle.com/uciml/sms-spam-collection-dataset>

⁷<https://github.com/hrwhisper/SpamMessage>

⁸<https://github.com/mathigatti/sexting-dataset>

⁹<http://www.midureader.com/>

¹⁰<https://scikit-learn.org/>

¹¹<https://github.com/fxsjy/jieba>

can constitute toxic texts according to the templates, since not all keyword pairs are suitable. For example, for the template of pornographic (*i.e.*, "A you/my B"), A should be a verb and B should be a sexual-related organ. It is inappropriate to fill a verb keyword into slot B, leading to meaningless sentences.

In order to find proper keyword pairs suitable for different templates, DUO has to first obtain the property of every keyword. Considering that a keyword may have multiple properties, for each keyword, DUO first retrieves 5 sentences containing the keyword from the corresponding dataset. Then, it utilizes the language analysis method to perform Part-of-Speech tagging (PoS tagging), which identifies the word property (*e.g.*, noun, verb, adjective or adverb), and Named Entity Recognition (NER), which determines whether the keyword is a belongs to a pre-defined category such as group names or location. The results of each keyword are voted on by the results of the five sentences. In the implementation, DUO adopts Flair toolkit¹² for English and Baidu NLP API¹³ for Chinese.

After obtaining the keyword property, DUO extracts keyword pairs for three types of toxic contents as follows:

For abuse and hate speech, the template is "A is/are B", where A is a group and B is a negative noun or adjective. Hence, for A DUO selects the keywords being identified as nouns by PoS tagging and group names by NER. And for B, DUO selects the keyword which is either a noun or adjective and also has a negative sentiment, which is obtained from the APIs provided by Google¹⁴ (for English) and Baidu¹⁵ (for Chinese). In addition, we require A and B to be different words, aiming to prevent the "monkey is monkey" situation that generates non-toxic seed sentences.

For malicious advertisement, the template is "A: B", where A is a product and B is the contact information. Hence, DUO selects the keyword identified as the noun for A, according to the PoS tagging toolkit. For B, DUO extracts the keywords with a prefix of a kind of contact way. We collect a candidate list, including "Tel", "Email", "WhatsApp" and "Ins". In this way, DUO can extract keyword pairs that can constitute the prefix contact way followed by the specific contact information, such as "Tel: 12345678".

For pornography, the template is "A your/my B", where A is a verb and B is a sexual-related organ. Hence, DUO extracts the keyword identified as verb for A and noun for B, according to the PoS tagging toolkit.

In this way, DUO can extract keyword pairs as seeds for the further test case (*i.e.*, multimedia content) generation. Particularly, DUO generates 100 keyword pairs for each type of toxicity and each language, ending up with $100 * 3 * 2 = 600$ keyword pairs.

3.2 Modality Transformation

Since the extracted keyword pairs only contain single modality information (*i.e.*, textual), in order to generate multi-modal contents, DUO has to perform modality transformation. Particularly, this work focuses on three kinds of modalities, *i.e.*, visual, textual, and audio. For each extracted keyword pair (A, B), DUO transforms the information of A and B to different modalities as follows:

Visual modality. For the keyword pair (A, B), DUO transforms the textual information to visual modality via calling the Google figure API to search the top 5 images with A or B as the query. To ensure that the returned figures correctly contain the information of the queried keyword, DUO further utilizes Baidu Image Recognition API¹⁶ to check the recognized salient object in each figure. If the object is not equal to the keyword, DUO discards the figure.

Textual modality. Since the keywords are texts *per se*, DUO directly uses the templates for keyword extraction (Section 3.1) to obtain the textual information.

Audio modality. In order to transform the textual information to audio modality data (*i.e.*, speech), DUO calls the Baidu text-to-speech synthesis API¹⁷.

3.3 Multi-modal Semantic Fusion

After obtaining the information of keyword pairs from different modalities, DUO performs semantic fusion, which fuses the multi-modal information together and generates test cases (*i.e.*, multimedia contents). In particular, DUO generates two kinds of multimedia data, *i.e.*, images and videos. The images are generated by fusing visual and textual information, while the videos can be constituted by any combination of the three modalities, *i.e.*, visual & textual, visual & audio, audio & text, and visual & textual & audio. Reminding that, besides the (A, B) keyword pair, there is a middle word in our template that represents the logical connection between A and B . For example, in abuse and hate speech, there is an "are" between A and B . We add this middle word to text or audio modality.

3.3.1 Image Generation. Image generation is performed via fusing visual and textual information (*i.e.*, image and text), more specifically, inserting the text into the image. In this process, we need to address three main issues: (1) how to decide the size of the text to avoid being too big or too small; (2) how to decide the location for the insertion, which should not affect the image content; (3) how to decide the color of the text so that it can be recognized in the image. In the following, we introduce three algorithms adopted by DUO to address these issues, respectively.

Text resizing. The goal of text resizing is to resize the text object to make the inserted text a comparable size with the salient object in the image. To achieve this target, DUO first utilizes Baidu Image Recognition API to detect the salient object in the image and obtains the coordinates of the four vertices of its bounding box $(x_1, y_1), (x_1, y_2), (x_2, y_1), (x_2, y_2)$. Then, the height h and the width w of the salient object can be approximated as $w = x_2 - x_1$ and $h = y_2 - y_1$, respectively. Finally, DUO sets the area of text within the range of $[0.8 * w * h, 1.2 * w * h]$, where 0.8 and 1.2 are hyper-parameters that we manually set based on empirical experiences.

Location determination. After resizing the text, DUO should select a suitable location in the image to insert the text. An ideal location should (1) have few overlaps with the salient object in the image, which is beneficial for humans to recognize both the text and the image; (2) the relative positions of the image and text follow the reading habits which is easier for human to understand the logic relation between the image information and the text information.

¹²<https://github.com/flairNLP/flair>

¹³https://ai.baidu.com/tech/nlp_basic/

¹⁴<https://cloud.google.com/natural-language>

¹⁵https://ai.baidu.com/tech/nlp_apply/sentiment_classify

¹⁶<https://ai.baidu.com/ai-doc/IMAGERECOGNITION/Xk3bcxdum>

¹⁷<https://ai.baidu.com/ai-doc/SPEECH/jk38y8gno>

To find the location with few overlaps with the salient object, we first define 9 candidate insert positions: top left, top middle, top right, middle left, middle middle, middle right, bottom left, bottom middle and bottom right. Since DUO has obtained the coordinates of the four vertices of the salient object in text resizing, it can directly calculate the overlapping area between the salient object and each candidate position. If the overlapping area is larger than 30%, DUO discards this candidate position.

To find the position in line with human reading habits, DUO utilizes a rule-based method based on the consideration that humans typically read from top to bottom and from left to right. It is worth noting that the keyword pair (A, B) is extracted according to the templates (*i.e.*, "A is/are B", "A: B", and "A your/my B"). In all the used templates, a human reads A first and reads B later. Hence, object A should be above or to the left of object B , such constraints help DUO filter out some candidate positions. For example, suppose we need to insert text B into image A , the candidate positions of B following human reading habits are middle right, bottom middle and bottom right. For each (image, text) pair, if there is no ideal location candidate left based on the criteria above, DUO discards this pair. If more than one candidates are suitable, DUO randomly selects one as the location to insert the text.

Font color selection. After determining a location to insert the text, another issue DUO needs to decide is the font color. If the font color is too similar to the background image color, the text will be hard for humans to recognize, leading to an invalid image. To mitigate this issue, DUO adopts a special webkit property called Text Stroke, which adds an exterior border around each character of the text. Text stroke can change the outline of the text, such as setting a color different from the original font color, so that the text can be recognized easily whatever the background color.

3.3.2 Video Generation. Unlike generating image data, video data can be generated by fusing any two or more information in different modalities. DUO conducts semantic fusion on different modality data as follows:

Fusing vision and text. Generating video test cases that fuse vision and text is similar to image generation, except that we should take the order of showing the image and the text into consideration. Similar to the position selection when generating images, for a keyword pair (A, B) , A should come first and B should come later. Hence DUO generates the video that presents A first and then presents B . For example, if A is an image and B is the text. The generated video shows A first for a while and then shows the text of B . On the other hand, if A is the text and B is an image, the video would show the text of A and then show the image B . In the implementation, DUO adopts `ffmpeg`¹⁸, a complete, cross-platform API to record, convert and stream audio and video.

Fusing vision and audio. DUO generates the video based on vision and audio by showing the image and playing the audio (a synthesized speech generated from the text). To make the video easier for humans to understand, again, we consider the order of showing the image and playing the audio. Specifically, for a keyword pair (A, B) , the generated video shows A first and then shows B . For example, if A is an image and B is the audio. We show A first for a while and then play the audio of B . On the other hand,

if A is the audio and B is an image, the video first plays the audio of A and meanwhile shows a blank video screen. The image B will not be displayed until audio A has finished playing.

Fusing text and audio. Generating the video test cases with audio and text information is similar to generating the video that fuses image and audio. The main difference is that here DUO shows the text, rather than the image, and plays audio. Again, if A is audio and B is text, the generated video plays the audio first and then shows the text, and *vice versa*. Since this kind of fusion contains both text and audio modality, we randomly add the middle word to either of the modality.

Finally, it is also feasible to generate videos by fusing all three modalities. To achieve this, DUO adopts the procedure of fusing image and audio. The only difference is that here DUO shows the middle words in text format between showing A and showing B , rather than showing accompanied with the keyword in audio modality.

As such, DUO can generate multi-modal images and videos, which are used as test cases to detect errors in content moderation software.

3.4 Toxicity Collection and Error Detection

After modality transformation and semantic fusion, DUO constructs test cases (*i.e.*, multi-modal images and videos). Each test case has a corresponding keyword pair (A, B) , which is extracted based on a specific template. By filling the keyword pair into the template, we can get a *seed sentence*, which is further used for solving the test oracle problem. During testing, DUO first feeds the generated test cases, as well as the corresponding seed sentences, to the content moderation software under test. Since all seed sentences are supposed to be toxic, any test case whose corresponding seed sentence is classified as non-toxic by the textual content moderation software will be discarded.

After toxicity collection, DUO inspects the predicted toxicity of each test case. Any case that violates the metamorphic relation (MR) will be reported as a suspicious error. The MR is designed based on the following simple consideration: the toxicity should remain the same when some of the information is transformed into another modality. Since DUO assumes seed sentences are all toxic ones, all the remaining cases should be classified as toxic, *i.e.*, an error is found if the test case is categorized as non-toxic.

In the implementation, we test 5 commercial software products provided by large Internet companies, *i.e.*, Google Cloud¹⁹, Amazon Web Service²⁰, Baidu Cloud²¹, Tencent Cloud²² and Alibaba Cloud²³, all of which are the official content moderation software from big technology companies with more than 100 millions of users. In particular, all the software products are the latest version by Nov. 1st, 2022, when the experiments were conducted. The version information of the software under test is listed in Table 2. Besides commercial software products, we also test popular research models. The core of multimedia content moderation software is

¹⁸<https://ffmpeg.org/>

¹⁹<https://cloud.google.com/video-intelligence/docs/analyze-safesearch>

²⁰<https://docs.aws.amazon.com/rekognition/latest/dg/moderation.html>

²¹<https://cloud.baidu.com/doc/ANTIPORN/s/6ki012lqu>

²²<https://cloud.tencent.com/document/product/1235>

²³https://help.aliyun.com/document_detail/146716.html

Table 2: Software Version Information.

Software	Version	Launch Date
Google	builtin/stable	2022.05.05
Amazon	5.0	2022.10.01
Baidu	4.16.3	2022.03.25
Tencent	2022-06-30	2022.06.30
Alibaba	2022.06.15	2022.06.15

image classification, hence we test Resnet-based image classification model²⁴ and Vision-Transformer-based image classification model²⁵, both having more than 100k downloads according to Hugging Face model zoo, a famous AI model repository.

4 EVALUATION

To validate the effectiveness of DUO and get more insights on enhancing content moderation software, we use our method to test five commercial software products and two state-of-the-art research models for multimedia content moderation. In this section, we detail the evaluation process and empirically explore the following four research questions (RQs).

- RQ1: Are the test cases generated by DUO toxic and realistic?
- RQ2: Can DUO find errors in content moderation software?
- RQ3: What factors affect the performance of DUO?
- RQ4: Can we use DUO to improve the performance of multimedia content moderation?

4.1 RQ1: Are the test cases generated by DUO be toxic and realistic?

DUO aims to generate test cases that are toxic and are as realistic as the ones real-world users produce to evade moderation. Thus, in this section, we evaluate whether the generated test cases are still toxic (i.e., semantic-preserving) and whether they are realistic.

We conduct human annotation via crowd-sourcing. First, we generate 10 images and 30 videos (10 vision + text, 10 vision + audio, 10 audio + text) for each task and each language, ending up with 240 test cases for annotation. For each test case, we ask three questions: 1) From "1 strongly disagree" to "5 strongly agree", to what extent do you agree that the image/video is semantically equivalent to the sentence? 2) From "1 strongly disagree" to "5 strongly agree", to what extent do you agree that the image/video is toxic (hate speech, malicious advertisement, or pornography, according to the dataset)? 3) From "1 strongly disagree" to "5 strongly agree", to what extent do you agree that this kind of image/video is realistic that Internet users would use? For English/Chinese, we distribute the questionnaire and recruit 20/20 crowd workers on Prolific²⁶/Tencent Wenjuan²⁷, who have English/Chinese as their first language. Before annotation, we provide instructions about the type of questions and asked them to make subjective judgments in the annotation. We do not provide additional training to avoid potential bias from us. Then the annotators are asked to annotate the test cases. Annotation results show that: 1) the generated test

²⁴<https://huggingface.co/microsoft/resnet-50>

²⁵<https://huggingface.co/google/vit-base-patch16-224>

²⁶<https://www.prolific.co/>

²⁷<https://wj.qq.com/>

cases are semantically equivalent to the seed sentence, with an average score of 4.46/4.59; 2) the generated test cases are toxic, with an average score of 4.19/4.29; 3) the generated test cases are realistic, with the average score of 3.96/4.15. We followed [32] to measure the inter-annotator agreement using Randolph's Kappa, obtaining a value of 0.84/0.81 for the test cases, which indicates "almost perfect agreement". There are a few cases that the generated images are annotated as non-toxic subjectively. For example, sometimes Google Image API returns a lovely cartoon image which could make the annotator feel less offensive.

Answer to RQ1: Test cases generated by DUO are toxic and commonly seen in real-world scenarios.

4.2 RQ2: Can DUO find errors in content moderation software?

DUO aims to automatically generate test cases to find potential bugs in current content moderation software. Hence, in this section, we evaluate the number of errors that DUO can find in the outputs of commercial content moderation software and academic models.

4.2.1 Software Products and Models under Test. We use DUO to test five commercial content moderation software products and two state-of-the-art research models.

Image Content Moderation. For commercial software products, we choose the products that have been deployed in the cloud of large Internet companies, including Google, Amazon, Baidu, Tencent, and Alibaba. All of them can be accessed by registered users via an API. In particular, since Amazon Web Service does not provide advertisement detection services, Google Cloud does not provide hateful image detection and advertisement detection services, and Tencent Cloud does not provide hateful and abusive image detection services, we do not include them in our experiments. For research models, we choose the Vision-Transformer Model and ResNet-18 Model, which are the most downloaded image classification models in Hugging Face²⁸ model zoo. For hate speech detection, we use the Hateful Memes Dataset [28], a hate and abuse detection dataset containing 8,500 multimodal memes with labeled toxicity. For malicious advertisement detection, we use the Advertisement Understanding Dataset [21], an advertisement dataset containing 64,832 Ad and 13,597 non-Ad images. Since there is no publicly available pornographic image detection dataset, we do not include this in our experiments.

We follow the official hyper-parameters setting to train the research model. More specifically, we fine-tuned the model with 15 epochs, a batch size of 8, and a weight decay of 0.01. We use the model with the highest accuracy on the validation set, and test its performance on the test set.

Video Content Moderation. All the five commercial software products mentioned above provide video content moderation services. Similar to the image content moderation experimental settings, we conduct experiments on the toxic content category if the APIs provide the corresponding detection services. Thus, there are no results

²⁸<https://huggingface.co/>

Table 3: Test case statistics.

Software	Tasks	Seed Num	Toxic Num
Google	Abuse&Hate	100	81
	pornography	100	74
	Advertisement	100	-
Baidu	Abuse&Hate	100	89
	pornography	100	95
	Advertisement	100	91
Tencent	Abuse&Hate	100	84
	pornography	100	89
	Advertisement	100	91
Alibaba	Abuse&Hate	100	74
	pornography	100	85
	Advertisement	100	84

on hate detection in Google Cloud and advertisement detection in Amazon Web Service, Google Cloud, and Tencent Cloud. Since the video content moderation service in Amazon Web Service and Google Cloud do not support audio modality analysis, we consider the modality fusion between vision and textual for these two APIs. Since there is no publicly available multi-modal video classification dataset for hate, pornographic, and advertisement detection, we do not conduct testing on the research model for video content moderation.

For each tested software product or model, we first input all the seed data and filter out the sentence that cannot be identified as toxic. In other words, we only used those that have already been recognized as toxic contents as seed data to generate test cases. The statistics of seed data are shown in Table 3, which shows that most of the generated seed sentences are toxic. Then, we generate image and video test cases accordingly. Finally, we use the generated test cases to test the software products and research models.

To evaluate how well DUO does on generating test cases that trigger errors, we calculate Error Finding Rate (EFR), which is defined as follows:

$$\text{EFR} = \frac{\text{the number of misclassified test cases}}{\text{the number of generated test cases}}$$

4.2.2 Analysis. We list the results in Table 4 and 5.

Overall Analysis. In general, DUO is able to find errors in software products and research models with a relatively high EFR. Since most of these test cases are annotated as toxic according to our annotators, we believe such high EFR implies the effectiveness of DUO and reveals the unexpected vulnerability of widely-deployed software products and research models.

One common concern about AI software testing is whether the software performs well on existing test cases, which are toxic single-modal inputs in content moderation. To address this concern, we conduct a lightweight experiment to evaluate the effectiveness of the software under test in detecting toxic contents from the Internet. Since there is no publicly-available toxic (hateful, porno, and malicious ad) image benchmark dataset (probably due to the toxic nature), we manually collect a dataset with 50 hateful images, 50

porno images, and 50 ad images from the Internet. The average detection rate of five content moderation software is 97.8%, indicating the effectiveness of the software. Thus, we think the high EFR achieved by DUO is exciting.

The comparison between image and video. The robustness of image content moderation and that of video content moderation do not strictly hold a positive correlation. For example, Baidu’s image moderation is much worse than its own video moderation, even though the testing logic should be similar. We think the algorithms behind them are different.

The comparison between different software. Since the seed sentence and test cases are different, it is difficult to compare the APIs for English content, *i.e.*, Amazon Web Service and Google Cloud, with the APIs for Chinese content, *i.e.*, Baidu Cloud, Tencent Cloud, and Alibaba Cloud. Even if the APIs for the same language are using the same test cases, it is also difficult to obtain an interesting conclusion because the functionalities of different software could vary a lot.

The comparison within a product. The performance of moderating different kinds of toxic data for one software product is different. For example, the robustness of pornography detection is much better than that of advertisement detection for the image content moderation APIs provided by Baidu, Tencent, and Alibaba. We think it is because the bulk of Internet Companies’ revenue comes from advertising so it is common to share the link, image, or video of advertisement between the users on these platforms. In addition, in Chinese culture, advertisement is not as toxic as pornography. Therefore, these companies seem to pay less attention to advertisement detection than pornography detection when developing their content moderation software.

Answer to RQ2: DUO can find substantial errors in both commercial software and research models.

4.3 RQ3: What factors affect the performance of DUO?

This section explores the impact of two factors on the DUO’s performance. First, we studied the impact of the location of the image and text. In theory, a video should be identified as toxic content if it contains a toxic image or text, no matter where the image or text is located. However, we found that changing the location of the image and text can affect the prediction of content moderation software. More specifically, we randomly select 20 seed sentences for each kind of toxicity and generate two kinds of vision + text video test cases with different locations: 1) the image and the text are showing on the left of the screen; 2) the image and the text are showing on the right of the screen. The test cases are fed into the video content moderation software products. We found that on average 3.2% of video test cases can bypass the moderation after we move the image and the text from the left to the right of the screen. This result implies that we can find the best location for toxic images and text if we want to find more cases that can bypass moderation.

Table 4: Error finding rates of image content moderation software and academic models (AM).

	Amazon	Google	AM_ResNet	AM_ViT	Baidu	Tencent	Alibaba
Abuse Detection	91.7%	-	94.1%	90.6%	99.55%	-	76.50%
Porno Detection	78.82%	77.18%	-	-	96.49%	76.77%	33.14%
Ad Detection	-	-	78.8%	75.3%	99.78%	99.29%	65.46%

Table 5: Error finding rates of video content moderation software.

Perturbation	Abuse Detection				Porn Detection					Ad Detection	
	Amazon	Baidu	Tencent	Alibaba	Amazon	Google	Baidu	Tencent	Alibaba	Baidu	Alibaba
V + T	89%	71.91%	100%	60.81%	91.38%	100%	50.53%	100%	60%	0%	100%
V + A	-	100%	100%	54.05%	-	-	91.58%	100%	55.29%	5.49%	100%
A + T	-	32.58%	100%	81.08%	-	-	10.53%	100%	58.82%	0%	100%
V + A + T	-	76.67%	100%	97.29%	-	-	55.79%	100%	97.64%	7.61%	100%

Second, we studied the impact of the selection of modality. In theory, a toxic content should be identified as toxic no matter which modality it is. For example, an abusive sentence should be categorized as toxic no matter if it is in text format or being converted to audio format. However, we found that for multimedia content moderation software, the sensitivity of different modalities is not the same. To this end, we randomly select 20 toxic seed sentences for each kind of toxicity and generate two kinds of video test cases: 1) pure text video that shows the sentence from the beginning to the end; 2) pure audio video that uses text-to-speech APIs to convert the seed sentence to audio and playing the audio. The test cases are fed into the video content moderation software products. We found that on average 94.7% of the pure text video can be detected as toxic while only 84.2% of the pure audio video can be detected. This result implies that we can distribute more toxic information to the less sensitive modality if we want to find more cases that can bypass moderation.

Answer to RQ3: The location of the image and text and the modality can affect the performance of DUO.

4.4 RQ4: Can we use DUO to improve the performance of multimedia content moderation?

In this section, we discuss how to use DUO to improve the robustness of multimedia content moderation. A natural way would be to retrain the models using the test cases generated by DUO and check whether the retrained models become more robust. This method is known as the robust retraining method. For each research model, we randomly select 100 test cases that successfully bypass moderation and label them as toxic. After that, we add these 100 labeled images to the corresponding training set and re-train the ResNet-18 model and the Vision-Transformer model. We adopt the default fine-tuning settings from the Hugging Face website. In other words, the setting of robust training is identical to the setting of normal training. The only difference is the data.

Table 6: Improvement of models with DUO

Model	Task	Ori	Rob
ResNet-18	Abuse	94.1%	5.7%
	Ad	78.8%	3.6%
ViT	Abuse	90.6%	4.8%
	Ad	75.3%	2.5%

To validate the effectiveness of robust retraining with DUO, we use DUO to test the model after robust retraining, denoted as "Rob", and compared its EFRs with the original model, denoted as "Ori". The results are presented in Table 6. We can observe that robust training with DUO's test cases largely reduces the EFRs on all the settings, which shows that DUO can effectively improve the robustness of multimedia content moderation models. In addition, we also evaluate their model accuracy on the original test set and the results show that in all settings, model accuracy remains the same after robust training.

Answer to RQ4: Test cases generated by DUO can effectively improve the robustness of multimedia content moderation models without affecting their accuracy.

5 DISCUSSION

5.1 Summary of Findings

5.1.1 Different Modalities Matter. In the evaluation, we can observe that different content moderation software has different specialties. For example, the content moderation software of Baidu performs quite well (*i.e.*, low EFR) on videos fused by audio and text, while cannot achieve satisfying performance on those fused by vision and audio. In addition, the software from Alibaba can achieve 65.46% EFR on image-based Ad detection, while it is infeasible to detect video-based advertisements (*i.e.*, 100% EFR for all kinds of fused videos). Such characteristics make it non-trivial to

comprehensively test content moderation software, *e.g.*, test cases for software may be inappropriate for another. In contrast, DUO considers three different modalities, and performs semantic fusion to generate both image and video test cases in English and Chinese. Such design makes DUO produce comprehensive test cases for different content moderation tasks. In addition, DUO is designed for black-box testing, which does not require white-box access to the software under test. Such design makes it convenient for adapting DUO to other AI-based software testing tasks, such as image caption systems.

5.1.2 Future Direction. Although DUO is able to find numerous errors in multimedia content moderation software and research models, there are several limitations that can lead to future works.

First, we can improve the diversity of our generated test cases. In this work, we only design one simple template and extract limited keywords from a toxic dataset for each kind of toxicity to generate test cases. In future work, we can follow the framework of this paper but use more kinds of templates and keywords.

Second, we can improve the authenticity of our generated test cases. In this work, we conduct multi-modal fusion by inserting text into an image or showing an image while playing a synthesized speech. In future work, we can utilize more advanced conditional image or video generation models to generate more vivid test cases.

5.2 Simplicity of the Methodology

While our *Semantic Fusion* methodology is simple, it is novel, general, and effective. The idea that fuses the semantics of single-modal inputs into a new input that combines the semantics of its ancestors has not been explored by any existing studies. Meanwhile, although the templates we use are simple, the single-model inputs can be automatically extracted from media sources or datasets to generate diverse test cases. It can be easily generalized to more complicated templates and seeds. In addition, the test cases generated from our templates can already reveal many errors in real-world software, which was reported to be highly effective by the companies. We believe that our simple methodology could benefit the community and we hope our work will be a starting point that brings more attention to multimedia content moderation in the field.

5.3 Threats to Validity

Internal Validity. The major threat to internal validity is that the image and video test cases generated by DUO may no longer preserve the toxic nature of the seed sentence, which may cause false positives during testing. Besides, DUO only uses simple templates to extract keyword pairs for test case generation. The template design can also affect the quality of the generated test cases, which may not be realistic enough that seems to be applied by users and appear in the real world. To relieve these issues, we validated the generated test cases by conducting human annotations, which shows that our generated text cases are toxic and realistic. During the manual inspection of the evaluation results, we also observe few false positive. In addition, DUO adopts existing tools for data transformation, *e.g.*, Flair for PoS tagging and ffmpeg for video processing, whose performance may affect the quality of generated cases. To relieve this issue, we only chose tools that are widely used by millions

of users. Moreover, these tools are also replaceable, which means users of DUO can customize these tools for specific testing tasks.

External Validity. The major threat to external validity is whether the content moderation software products we test are good and representative of what the industry uses. To relieve this issue, we evaluated the effectiveness of these products. The average detection rate of five content moderation software products is 97.8%, indicating their effectiveness. As for representativeness, all five commercial software products are paid cloud services provided by the cloud platform from the big companies. Moderation services and other cloud services have become an important source of income for them. Meanwhile, a huge amount of downstream companies and users are using paid services provided by these companies. The customer list can be seen from Google²⁹, Amazon³⁰ and Baidu³¹. Thus, we believe the moderation software studied in our paper has a significant impact on real users and is representative of industry practice. The other possible threat to external validity lies in the source datasets used for keyword extraction. Low-quality keywords can directly affect the testing effectiveness of generated test cases. To relieve this threat, we extracted keywords from six practical datasets, which consist of real-world toxic contents. Manual inspection and the evaluation results also prove that DUO can produce effective multimedia contents. Another threat may lie in the comprehensiveness and representation of our metamorphic relations, which might hurt the generalizability of our results and findings. To mitigate this threat, we extensively implemented splitting and fusion for all the possible permutations of different modalities. Hence, the final metamorphic relations are empirically comprehensive. To further validate the generalizability of DUO, we also conducted extensive experiments. Specifically, we tested DUO on five commercial textual software products and two research methods for content moderation. We also tested DUO with different kinds of toxic data: abuse hate speech, porn content and malicious advertisement, in both English and Chinese. The evaluation results confirm that our findings can generalize to different methods and tasks.

6 RELATED WORK

6.1 Content Moderation

There are generally two categories of research models for content moderation, *i.e.*, feature engineering-based models and neural networks-based models. Feature engineering-based models work by manually designing features and using the features to train classification models to identify toxic contents. While a neural networks-based method is typically a neural network model trained in an end-to-end manner on a huge amount of data, which does not require much human effort on feature engineering.

6.1.1 Single-modal Content Moderation. In the literature, there are mainly three categories of data modalities, *i.e.*, visual, audio, and textual.

Visual Content Moderation. Early works on toxic image detection mainly extract some pre-designed features and then train

²⁹https://cloud.google.com/customers#/products=Data_Analytics

³⁰<https://aws.amazon.com/machine-learning/customers/>

³¹<https://cloud.baidu.com/partner/plan.html#search>

a machine learning classifier. For example, Shen *et al.* [54] computed color histograms to detect pornography images, depending on the activity suspicion that pixels in pornographic images are mostly skin. Zhang *et al.* [78] extract the motion features and use a Hidden Markov Model (HMM) for video anomaly detection. With the development of deep learning techniques and the building of big datasets, the progress of image representation learning has motivated researchers to explore neural network-based models for visual content moderation. For example, Moustafa [42] adapted a CNN architecture for image classification to the pornographic video classification task. Each frame is input for being classified as porn or non-porn and then integrating the final result for a video via a majority voting process.

Audio Content Moderation. Zhang [78] extracted Mel-Frequency Cepstral Coefficients (MFCCs) feature and use an HMM to detect audio events. [56] extract various audio features, including MFCC, mean of short-time zero-crossing rates (mSTZCR), mean of the spectral centroid (SC) and high short-time zero-crossing rates ratio (HZCRR), and use in-class clustering for porno-sounds detection. Gupta [16] utilized a bidirectional Gated Recurrent Units (GRU) and Long Short-Term Memory (LSTM) model to detect abuse from multilingual audio. Recently researchers proposed a deep learning method that can detect COVID-19 from breath and cough audio [7].

Textual Content Moderation. Early works are rule-based methods that are based on pre-defined rules or dictionaries [45, 59]. However, rule-based can hardly deal with implicit abuse and sarcasm. Besides, they are vulnerable to errors in spelling, punctuation, and grammar [70]. Computation-based methods leverage some statistics of the textual data, such as TF-IDF [51, 74]. Computation-based methods require less human effort and are more robust to spelling, punctuation, and grammatical variations. Nevertheless, this kind of approach can only capture surface-level patterns, not deeper semantic properties [70]. The advancements in text representation learning have motivated researchers to explore neural network-based models, such as Feed-forward neural networks [11], LSTM [1] or the pre-training models [10, 37], on textual content moderation task and have achieved remarkable performance.

6.1.2 Testing Multi-modal Content Moderation. The testing of multi-modal content moderation is relatively unexplored compared with that of single-modal content moderation. Kiela *et al.* [28] manually created and labeled a multi-modal hateful memes dataset as a benchmark to evaluate the multi-modal understanding ability of cross-vision and text modality. Hussain *et al.* [21] collected and labeled a TV advertisement classification benchmark that predicts whether a frame of video is an advertisement. DUO is different from the existing papers: 1) existing work needs huge human effort to create and label the multi-modal test cases while DUO can automatically generate and label the test cases; 2) existing work only provides image test cases on English hateful and advertisement image detection, while DUO is able to comprehensively generate image and video test cases in two languages for three kinds of toxicity detection. In addition, our experiments show that DUO can find numerous errors from models trained on the datasets provided by existing work. We believe our work complements this line of existing papers.

6.2 Robustness of AI Software

Artificial Intelligence (AI) software has been widely adopted by many domains, such as autonomous driving and face recognition. However, AI software is not robust enough and can generate erroneous outputs that lead to fatal accidents [34, 83]. To this end, researchers have proposed a variety of methods to generate adversarial examples or test cases that can fool AI software [4, 5, 13, 23, 27, 29, 30, 38–40, 44, 50, 60–62, 64, 68, 69, 79–81].

AI software has also become adept at solving many NLP tasks, such as sentiment analysis [82], reading comprehension [75], Grammatical Error Correction [72] and machine translation [2, 24, 25]. In recent years, inspired by the work on adversarial examples in the AI field, NLP researchers have started exploring attack and defense techniques for various NLP software [15, 19, 63]. For example, Ribeiro *et al.* [46] designed a behavioral testing method to test NLP software for sentiment analysis, duplicate question answering, and machine comprehension. Li *et al.* [35] used deep learning models to generate test cases for deep learning-based NLP software. And Wang *et al.* [67] proposed a metamorphic testing framework for textual content Moderation software. Unlike these studies, this work focuses on the robustness of multimedia content moderation software, which has not been explored by existing work.

7 CONCLUSION

In this paper, we propose a general, effective methodology, *Semantic Fusion*, for validating multimedia content moderation software and further realize it as a practical tool DUO. DUO can generate image and video test cases covering three kinds of toxic contents, including abusive language and hate speech, malicious advertisements, and pornography, in both two language settings (English and Chinese). We used DUO to test five commercial content moderation software products and two research models. Results show that DUO can find numerous errors from all the software and models. In addition, we showed that we can improve the models' robustness by utilizing the test cases produced by DUO to perform robust training. We hope our framework can help identify the defects of current content moderation software and facilitate their development, contributing to a cleaner and better Internet environment.

ACKNOWLEDGEMENT

The work described in this paper was supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (No. CUHK 14206921 of the General Research Fund) and the National Natural Science Foundation of China (Grant Nos. 62102340 and 62206318)

REFERENCES

- [1] Pinkesh Badjatiya, Shashank Gupta, Manish Gupta, and Vasudeva Varma. 2017. Deep Learning for Hate Speech Detection in Tweets. *Proceedings of the 26th International Conference on World Wide Web Companion* (2017).
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural Machine Translation by Jointly Learning to Align and Translate. *ICLR abs/1409.0473* (2015).
- [3] Helen Brown. 2022. The surprising power of internet memes. <https://www.bbc.com/future/article/20220928-the-surprising-power-of-internet-memes>. Accessed: 2022-9-01.
- [4] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuanhai Zhang, Michael E. Sherr, Clay Shields, David A. Wagner, and Wencho Zhou. 2016. Hidden Voice Commands. In *USENIX Security Symposium*.

- [5] Junjie Chen, Zhuo Wu, Zan Wang, Hanmo You, Lingming Zhang, and Ming Yan. 2020. Practical Accuracy Estimation for Efficient Deep Neural Network Testing. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 29 (2020), 1–35.
- [6] Yang Chen, Rongfeng Zheng, Anmin Zhou, Shan Liao, and Liang Liu. 2020. Automatic Detection of Pornographic and Gambling Websites Based on Visual and Textual Content Using a Decision Mechanism. *Sensors (Basel, Switzerland)* 20 (2020).
- [7] Harry Coppock, Alexander Gaskell, Panagiotis Tzirakis, Alice Baird, Lyn Jones, and Björn W. Schuller. 2021. End-to-end convolutional neural network enables COVID-19 detection from breath and cough audio: a pilot study. *BMJ Innovations* 7 (2021), 356–362.
- [8] Nikolina Cveticanin. 2022. What's On the Other Side of Your Inbox - 20 SPAM Statistics for 2022. <https://dataprof.net/statistics/spam-statistics/>. Accessed: 2022-03-01.
- [9] Statista Research Department. 2022. Latin America: adult online content accessed by children 2017-2020. <https://www.statista.com/statistics/1060619/share-adult-content-websites-visit-children-latin-america/>. Accessed: 2022-03-01.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *NAACL abs/1810.04805* (2019).
- [11] Nemanja Djuric, Jing Zhou, Robin Morris, Mihajlo Grbovic, Vladan Radosavljevic, and Narayan L. Bhamidipati. 2015. Hate Speech Detection with Comment Embeddings. *Proceedings of the 24th International Conference on World Wide Web* (2015).
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR* (2021).
- [13] Xinyu Gao, Yang Feng, Yining Yin, Zixi Liu, Zhenyu Chen, and Baowen Xu. 2022. Adaptive Test Selection for Deep Neural Networks. *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE) (2022)*, 73–85.
- [14] Tarleton Gillespie. 2020. Content moderation, AI, and the question of scale. *Big Data & Society* 7, 2 (2020), 2053951720943234.
- [15] Shashij Gupta. 2020. Machine Translation Testing via Pathological Invariance. *2020 IEEE/ACM 42nd International Conference on Software Engineering: Companion Proceedings (ICSE-Companion) (2020)*, 107–109.
- [16] Vikram Gupta, Rimi A. Sharon, Ramit Sawhney, and Debdeep Mukherjee. 2022. ADIMA: Abuse Detection In Multilingual Audio. In *ICASSP*.
- [17] Laura Hanu, James Thewlis, and Sasha Haco. 2021. How AI Is Learning to Identify Toxic Online Content. <https://www.scientificamerican.com/article/can-ai-identify-toxic-online-content/>. Accessed: 2022-03-01.
- [18] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016), 770–778.
- [19] Pinjia He, Clara Meister, and Zhendong Su. 2021. Testing Machine Translation via Referential Transparency. *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE) (2021)*, 410–422.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* (1997).
- [21] Zaeem Hussain, Mingda Zhang, Xiaozhong Zhang, Keren Ye, Christopher Thomas, Zuha Agha, Nathan Ong, and Adriana Kovashka. 2017. Automatic Understanding of Image and Video Advertisements. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017).
- [22] Michael Ibañez, Ranz Sapinit, Lloyd Lois Antonie Reyes, Mohammed Hussien, Joseph Marvin Imperial, and Ramon Rodriguez. 2021. Audio-Based Hate Speech Classification from Online Short-Form Videos. *2021 International Conference on Asian Language Processing (IALP)* (2021).
- [23] Pin Ji, Yang Feng, Jia Liu, Zhihong Zhao, and Baowen Xu. 2021. Automated Testing for Machine Translation via Constituency Invariance. *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE) (2021)*, 468–479.
- [24] Wenxiang Jiao, Zhaopeng Tu, Jiarui Li, Wenxuan Wang, Jen tse Huang, and Shuming Shi. 2022. Tencent's Multilingual Machine Translation System for WMT22 Large-Scale African Languages. In *Conference on Machine Translation*.
- [25] Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, and Zhaopeng Tu. 2023. Is ChatGPT A Good Translator? A Preliminary Study. *ArXiv abs/2301.08745* (2023).
- [26] Meng Jing. 2018. China's Baidu turns to AI to police online content, but is the technology reliable? https://www.scmp.com/tech/innovation/article/2143759/chinas-baidu-turns-ai-police-online-content-technology-reliable?module=perpetual_scroll_0&pgtype=article&campaign=2143759. Accessed: 2022-03-01.
- [27] Sungmin Kang, Robert Feldt, and Shin Yoo. 2020. SINVAD: Search-based Image Space Navigation for DNN Image Classifier Test Input Generation. *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops* (2020).
- [28] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes. *NeurIPS* (2020).
- [29] Junhwi Kim, Minhyuk Kwon, and Shin Yoo. 2018. Generating Test Input with Deep Reinforcement Learning. *2018 IEEE/ACM 11th International Workshop on Search-Based Software Testing (SBST)* (2018), 51–58.
- [30] Seah Kim and Shin Yoo. 2021. Multimodal Surprise Adequacy Analysis of Inputs for Natural Language Processing DNN Models. *2021 IEEE/ACM International Conference on Automation of Software Test (AST)* (2021), 80–89.
- [31] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *EMNLP*.
- [32] Hannah Rose Kirk, Bertram Vidgen, Paul Röttger, Tristan Thrush, and Scott A. Hale. 2021. Hatemoji: A Test Suite and Adversarially-Generated Dataset for Benchmarking and Detecting Emoji-based Hate. *ACL abs/2108.05921* (2021).
- [33] Merlene Leano. 2020. 5 Examples of Content Moderation Failures that Prove Why Content Moderation is an Essential Tool. <https://www.linkedin.com/pulse/5-examples-content-moderation-failures-prove-why-essential-leano/>. Accessed: 2022-08-01.
- [34] Sam Levin. 2018. Tesla fatal crash: 'autopilot' mode sped up car before driver killed, report finds [Online]. <https://www.theguardian.com/technology/2018/jun/07/tesla-fatal-crash-silicon-valley-autopilot-mode-report>. Accessed: 2018-06.
- [35] Linyang Li, Ruotian Ma, Qipeng Guo, X. Xue, and Xipeng Qiu. 2020. BERT-ATTACK: Adversarial Attack against BERT Using BERT. *EMNLP abs/2004.09984* (2020).
- [36] Zhou Li, Kehuan Zhang, Yinglian Xie, Fang Yu, and Xiaofeng Wang. 2012. Knowing your enemy: understanding and detecting malicious web advertising. *Proceedings of the 2012 ACM conference on Computer and communications security* (2012).
- [37] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *ArXiv abs/1907.11692* (2019).
- [38] Zixi Liu, Yang Feng, and Zhenyu Chen. 2021. DialTest: automated testing for recurrent-neural-network-driven dialogue systems. *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis* (2021).
- [39] Zixi Liu, Yang Feng, Yining Yin, and Zhenyu Chen. 2022. DeepState: Selecting Test Suites to Enhance the Robustness of Recurrent Neural Networks. *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE) (2022)*, 598–609.
- [40] Yuanfu Luo, Malika Meghiani, Qi Heng Ho, David Hsu, and Daniela Rus. 2021. Interactive Planning for Autonomous Urban Driving in Adversarial Scenarios. *2021 IEEE International Conference on Robotics and Automation (ICRA) (2021)*, 5261–5267.
- [41] Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2019. Tackling Online Abuse: A Survey of Automated Abuse Detection Methods. *ArXiv abs/1908.06024* (2019).
- [42] Mohamed N. Moustafa. 2015. Applying deep learning to classify pornographic images and videos. *PSIVT abs/1511.08899* (2015).
- [43] Emily R Munro. 2011. The protection of children online: a brief scoping review to identify vulnerable groups. *Childhood Wellbeing Research Centre* (2011).
- [44] Jie Zhang Haoyu Wang Shuang Liu Menghan Tian Qingchao Shen, Junjie Chen. 2022. Natural Test Generation for Precise Testing of Question Answering Software. *The 37th IEEE/ACM International Conference on Automated Software Engineering* (2022).
- [45] Amir H. Razavi, Diana Inkpen, Sasha Uritsky, and Stan Matwin. 2010. Offensive Language Detection Using Multi-level Classification. In *Advances in Artificial Intelligence*.
- [46] Marco Tulio Ribeiro, Tongshuang Sherry Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. In *ACL*.
- [47] Paul Röttger, Haitham Seelawi, Debora Nozza, Zeerak Talat, and Bertie Vidgen. 2022. Multilingual HateCheck: Functional Tests for Multilingual Hate Speech Detection Models. *NAACL* (2022).
- [48] Paul Röttger, Bertram Vidgen, Dong Nguyen, Zeerak Waseem, Helen Z. Margetts, and Janet B. Pierrehumbert. 2021. HateCheck: Functional Tests for Hate Speech Detection Models. In *ACL/IJCNLP*.
- [49] Henry A. Rowley, Yushi Jing, and Shumet Baluja. 2006. Large scale image-based adult-content filtering. In *VISAPP*.
- [50] Baptiste Rozière, J Zhang, François Charton, Mark Harman, Gabriel Synnaeve, and Guillaume Lample. 2022. Leveraging Automated Unit Tests for Unsupervised Code Translation. *ICLR* (2022).
- [51] Joni O. Salminen, Hind Almerexhi, Milica Milenkovic, Soon-Gyo Jung, Jisun An, Haewoon Kwak, and Bernard Jim Jansen. 2018. Anatomy of Online Hate: Developing a Taxonomy and Machine Learning Models for Identifying and Classifying Hate in Online News Media. In *ICWSM*.
- [52] Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. Social Bias Frames: Reasoning about Social and Power Implications of Language. *ACL* (2020).
- [53] Anna Schmidt and Michael Wiegand. 2017. A Survey on Hate Speech Detection using Natural Language Processing. In *SocialNLP@EACL*.

- [54] Xuanjing Shen, Wei Wei, and Qingji Qian. 2010. A pornographic image filtering model based on erotic part. *2010 3rd International Congress on Image and Signal Processing 5* (2010), 2473–2477.
- [55] Jack Shepherd. 2022. 30 Vital Video Marketing Statistics You Need to Know in 2022. <https://thesocialshepherd.com/blog/video-marketing-statistics>. Accessed: 2022-08-01.
- [56] Ziqiang Shi, Boyang Gao, Tieran Zheng, and Jiqing Han. 2009. Objectionable audio content recognition based on in-Class Clustering method. *2009 IEEE International Conference on Network Infrastructure and Digital Content* (2009), 712–716.
- [57] Adam Smith. 2019. New Zealand Terrorist Attack Live Streamed on Facebook. <https://www.pcmag.com/news/new-zealand-terrorist-attack-live-streamed-on-facebook>. Accessed: 2022-08-01.
- [58] Kaisong Song, Yangyang Kang, Wei Gao, Zhe Gao, Changlong Sun, and Xiaozhong Liu. 2021. Evidence Aware Neural Pornographic Text Identification for Child Protection. In *AAAI*.
- [59] Ellen Spertus. 1997. Smokey: Automatic Recognition of Hostile Messages. In *AAAI/IAAI*.
- [60] Zeyu Sun, Jie M. Zhang, Mark Harman, Mike Papadakis, and Lu Zhang. 2020. Automatic Testing and Improvement of Machine Translation. *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)* (2020), 974–985.
- [61] Yuchi Tian, Kexin Pei, Suman Sekhar Jana, and Baishakhi Ray. 2018. DeepTest: Automated Testing of Deep-Neural-Network-Driven Autonomous Cars. *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)* (2018), 303–314.
- [62] Yuchi Tian, Ziyuan Zhong, Vicente Ordóñez, Gail E. Kaiser, and Baishakhi Ray. 2020. Testing DNN Image Classifiers for Confusion & Bias Errors. *2020 IEEE/ACM 42nd International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)* (2020), 304–305.
- [63] Jen tse Huang, Jianping Zhang, Wenxuan Wang, Pinjia He, Yuxin Su, and Michael R. Lyu. 2022. AEON: a method for automatic evaluation of NLP test cases. *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis* (2022).
- [64] James Tu, Huichen Li, Xinchun Yan, Mengye Ren, Yun Chen, Ming Liang, Eilyan Bitar, Ersin Yumer, and Raquel Urtasun. 2021. Exploring Adversarial Robustness of Multi-Sensor Perception Systems in Self Driving. *ArXiv abs/2101.06784* (2021).
- [65] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. *NeurIPS* (2017).
- [66] James Vincent. 2020. Facebook is now using AI to sort content for quicker moderation. <https://www.theverge.com/2020/11/13/21562596/facebook-ai-moderation>. Accessed: 2022-03-01.
- [67] Wenxuan Wang, Jen tse Huang, Weibin Wu, Jianping Zhang, Yizhan Huang, Shuqing Li, Pinjia He, and Michael R. Lyu. 2023. MTM: Metamorphic Testing for Textual Content Moderation Software. *ArXiv abs/2302.05706* (2023).
- [68] Zan Wang, Hanmo You, Junjie Chen, Yingyi Zhang, Xuyuan Dong, and Wenbin Zhang. 2021. Prioritizing Test Inputs for Deep Neural Networks via Mutation Analysis. *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)* (2021), 397–409.
- [69] Anjiang Wei, Y. Deng, Chenyuan Yang, and Lingming Zhang. 2022. Free Lunch for Testing: Fuzzing Deep-Learning Libraries from Open Source. *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)* (2022), 995–1007.
- [70] Michael Wiegand, Josef Ruppenhofer, Anna Schmidt, and Clayton Greenberg. 2018. Inducing a Lexicon of Abusive Words – a Feature-Based Approach. In *NAACL*.
- [71] Dominik Winterer, Chengyu Zhang, and Zhendong Su. 2020. Validating SMT solvers via semantic fusion. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. 718–730.
- [72] Hao Wu, Wenxuan Wang, Yuxuan Wan, Wenxiang Jiao, and Michael R. Lyu. 2023. ChatGPT or Grammarly? Evaluating ChatGPT on Grammatical Error Correction Benchmark. *ArXiv abs/2303.13648* (2023).
- [73] Tingmin Wu, Sheng Wen, Yang Xiang, and Wanlei Zhou. 2018. Twitter spam detection: Survey of new approaches and comparative study. *Comput. Secur.* 76 (2018), 265–284.
- [74] Dawei Yin, Zhenzhen Xue, Liangjie Hong, Brian D. Davison, and Lynne Edwards. 2009. Detection of Harassment on Web 2.0.
- [75] Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohamad Norouzi, and Quoc V. Le. 2018. QANet: Combining Local Convolution with Global Self-Attention for Reading Comprehension. *ICLR abs/1804.09541* (2018).
- [76] Tai-Kuei Yu and Cheng-Min Chao. 2016. Internet Misconduct Impact Adolescent Mental Health in Taiwan: The Moderating Roles of Internet Addiction. *International Journal of Mental Health and Addiction* 14 (2016), 921–936.
- [77] A. A. Zaidan, Hezrul Abdul Karim, Nurul Nadia Ahmad, Bilal Bahaa Zaidan, and Miss Laiha Mat Kiah. 2015. Robust Pornography Classification Solving the Image Size Variation Problem Based on Multi-Agent Learning. *J. Circuits Syst. Comput.* 24 (2015), 1550023:1–1550023:37.
- [78] Dong Zhang, Daniel Gática-Pérez, Samy Bengio, and Iain McCowan. 2005. Semi-supervised adapted HMMs for unusual event detection. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)* 1 (2005), 611–618 vol. 1.
- [79] J Zhang, Mark Harman, Lei Ma, and Yang Liu. 2022. Machine Learning Testing: Survey, Landscapes and Horizons. *IEEE Transactions on Software Engineering* 48 (2022), 1–36.
- [80] Jianping Zhang, Jen tse Huang, Wenxuan Wang, Yichen Li, Weibin Wu, Xiaosen Wang, Yuxin Su, and Michael R. Lyu. 2023. Improving the Transferability of Adversarial Samples by Path-Augmented Method. *ArXiv abs/2303.15735* (2023).
- [81] Jianping Zhang, Weibin Wu, Jen tse Huang, Yizhan Huang, Wenxuan Wang, Yuxin Su, and Michael R. Lyu. 2022. Improving Adversarial Transferability via Neuron Attribution-based Attacks. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022), 14973–14982.
- [82] Lei Zhang and B. Liu. 2017. Sentiment Analysis and Opinion Mining. In *Encyclopedia of Machine Learning and Data Mining*.
- [83] Chris Ziegler. 2016. A google self-driving car caused a crash for the first time. [Online]. <https://www.theverge.com/2016/2/29/11134344/google-self-driving-car-crash-report>. Accessed: 2016-09.

Received 2023-02-16; accepted 2023-05-03