

Enhanced Multimodal Content Moderation of Children's Videos using Audiovisual Fusion

Syed Hammad Ahmed Muhammad Junaid Khan Gita Sukthankar

Department of Computer Science, University of Central Florida, Orlando, FL USA
{syed.hammad.ahmed, muhammad.junaid.khan, gita.sukthankar}@ucf.edu

Abstract

Due to the rise in video content creation targeted towards children, there is a need for robust content moderation schemes for video hosting platforms. A video that is visually benign may include audio content that is inappropriate for young children while being impossible to detect with a unimodal content moderation system. Popular video hosting platforms for children such as YouTube Kids still publish videos which contain audio content that is not conducive to a child's healthy behavioral and physical development. A robust classification of malicious videos requires audio representations in addition to video features. However, recent content moderation approaches rarely employ multimodal architectures that explicitly consider non-speech audio cues. To address this, we present an efficient adaptation of CLIP (Contrastive Language-Image Pre-training) that can leverage contextual audio cues for enhanced content moderation. We incorporate 1) the audio modality and 2) prompt learning, while keeping the backbone modules of each modality frozen. We conduct our experiments on a multimodal version of the MOB (Malicious or Benign) dataset in supervised and few-shot settings.

Introduction

There has been a remarkable surge in both the creation and consumption of video content on the internet, making automated content moderation indispensable for video sharing platforms. More than 80% of the internet usage globally is from video data streaming and downloading (Shewale 2024). In many countries, content moderation is mandatory for video hosting platforms to remain compliant with legal regulations. Recent findings indicate that around 2,500 videos are published per minute on YouTube, which is the largest video sharing platform (Hayes 2024). Unfortunately, there is a significant amount of content which can be detrimental to emotional and psychological well-being of the viewer. Extensive consumption of malicious video content has socio-economic impacts in addition to its effects on individual viewers.

Lawsuits have been filed against various online content sharing platforms for intentionally fostering addiction among children through their content, thereby exacerbating

the mental health crisis among the youth (Ortutay 2023). Video sharing websites combine both manual and automated content moderation approaches to eliminate the presence of published malicious content (YouTube 2024; TikTok 2024). For this paper, we sampled the content currently available for pre-schoolers on YouTube Kids and identified a number of malicious videos that haven't been flagged by the existing content moderation system. Figures 1 and 2 show snapshots from two different problematic videos which include violence and disturbing music, respectively. These examples highlight the need to develop robust and efficient tools for video content moderation.

Young children, particularly toddlers and pre-schoolers, are the demographic most susceptible to being adversely affected by video content created with malicious intent. Firstly, they are too naive to appreciate that certain audiovisual features are deliberately incorporated in videos in order to lure them to continue watching. Furthermore, continuous adult supervision is not pragmatic - in fact child's screen time serves as a babysitting tool, enabling parents to attend to other responsibilities or take a break, as reported by 32% of parents (Hinkley and McCann 2018). Numerous psychological studies discuss the adverse effects of exposure of inappropriate video content on a young child's mental and behavioral health.

Ahmed et al. (2023) provide a catalog of features that characterize malicious video content including quick and repetitive movements, frightening or repulsive appearance, harmful or destructive actions, and inappropriate or indecent behavior. They also note that malicious audio may include high-volume music or noise, cries or shouts, sounds of explosions or gunshots, and the use of offensive language. Excessive exposure to these inappropriate videos poses a notable threat to the cognitive growth of preschoolers.

Recent works on content moderation for children's videos have attempted to address the problem with multimodal approaches. Chuttur and Nazurally (2022) analyze text in the form of user comments and captions, along with image data. Their method is unable to capture the spatio-temporal context in videos as they use single images as inputs. Le, Tandon, and Oinar (2022) proposed a multimodal solution that incorporates subtitles and video metadata as input. As the subtitles do not contain information about music or sound effects, Samba may classify a video with malicious audio

content as safe. Papadamou et al. (2021) rely on video meta-data such as tags, title, statistics, in addition to video thumbnails. None of these models include audio inputs which is an essential modality for devising a robust video content moderation system. Tahir et al. (2020) and Alghowinem (2019) combined audio spectrograms with video features to classify the video content being either appropriate or inappropriate. However, the former focuses on detecting fake cartoon videos, whereas the latter only introduces a high-level idea without implementation.

Our proposed work bridges the gap in previous works by adapting OpenAI’s CLIP architecture (Radford et al. 2021) for videos, and enabling it to learn prompts for both the vision and text branches. Furthermore, we incorporate a pre-trained audio encoder from AudioCLIP (Guzhov et al. 2021b) to handle the audio modality while adding a fully-trainable projection layer to improve performance on the downstream task of detecting malicious videos. We also introduce a multimodal dataset with both audio and video modalities that includes annotations specifically targeted for the problem of content moderation for children’s videos.

Psychological research indicates that prolonged exposure to animated content featuring rapid movements can result in deterioration in performing everyday tasks (Lillard and Peterson 2011). Likewise, the impact of loud noises on a child’s brain development, particularly affecting reading, writing, and comprehension skills, has been established (Klatte, Bergström, and Lachmann 2013). The direct influence of violent cartoons on the behavior of preschoolers is evident, leading to increased levels of aggression and anxiety (Hapkiwicz 1979). Based on these psychological studies, in this paper we identify subtly harmful acoustic and visual features which makes the content alluring and addictive to a young child. These include a) brightly-hued fast moving objects with fast and loud music and b) split-screen animations with fast and loud music. A sample video snapshot for a) is shown in Figure 2.

This paper makes the following contributions towards the problem of content moderation of children’s videos:

1. Introduces a multimodal framework that includes the audio modality for more robust content moderation of children’s cartoon videos. We perform ablations to demonstrate the enhanced performance of adding audio.
2. Adapts CLIP to enable learning of prompts in both vision and text branches across multiple layers of encoders, along with AudioCLIP’s pre-trained audio encoder. We integrate a learnable projection layer in order to efficiently learn audio representations for the downstream content moderation task, while keeping the weights in the core models frozen.
3. Presents a comprehensive multimodal dataset MMOB (Multimodal Malicious or Benign) which includes labeled cartoon video samples with annotations of vision and audio labels. Baseline results for MMOB in supervised and few-shot learning settings are also showcased. We have released the dataset to encourage further research in the area.



Figure 1: One of the malicious video examples currently available on the YouTube Kids platform that shows a furious cartoon character shooting at other cartoon characters with a machine gun.

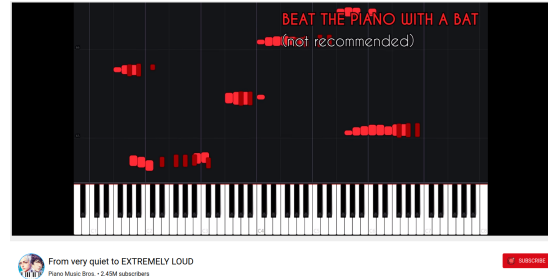


Figure 2: A sample video which includes fast and loud piano notes. Usually in such videos, there is high tempo, a lack of rhythm and varying pitches. The video also includes bright and striking hues, and suggests the violent action of “hitting the piano with a bat”.

Background

COPPA - Children’s Online Privacy Protection Act

The Children’s Online Privacy Protection Act (COPPA) (COPPA 1998), enforced by the US Federal Trade Commission (FTC), establishes clear privacy and content guidelines for online content publishers targeting children aged 13 or younger. In 2019, additional provisions were incorporated into COPPA, particularly addressing YouTube, to guarantee the delivery of safe content for children. In January 2024, the Federal Trade Commission (FTC) released a proposal revising the regulations that enforce the Children’s Online Privacy Protection Act (COPPA Rule). The proposed modifications aim to adapt to technological changes, and strengthen protections for children’s personal information based on the FTC’s review of public comments and enforcement experience (FTC 2024). Therefore, the problem of publishing age-appropriate content that we address in this paper, has been given high importance by FTC, which recommends avoiding content that could be harmful or inappropriate for children.

CLIP - Contrastive Language-Image Pre-Training

CLIP (Radford et al. 2021) is a very large multimodal model trained on a huge corpus of (image, text) pairs. Paired images and text are used for training both image and text encoders, jointly. The user provides the CLIP model with the text prompt that will elicit the best image classification; this text may include ground truth class information.

While learning the combined multimodal embedding space the model aims to maximize the cosine similarity scores of the image with the ground truth text while minimizing the cosine similarity between the embeddings of incorrect pairings. CLIP is capable of performing zero-shot classification and other downstream computer vision tasks efficiently using natural language supervision. In this work we modify CLIP for video input, and augment it to learn prompts in both the visual and textual branches of CLIP.

Related Work

Audio Modality

Audio is an oft neglected modality; however we believe that raw audio provides valuable contextual clues that are not included in the text transcription. This is particularly true in children’s cartoons where sound effects are often used in lieu of words to convey meaning to young children with smaller vocabularies. Popular cartoon sound effects include springs, tires screeching, slide whistles, chattering teeth, explosions, suction cups and rattles. Previous work has looked explicitly at the problem of sound event classification (McLoughlin et al. 2015), which could be potentially used to detect violent actions in children’s cartoons. A subcategory of this research focuses solely on environmental sound classification (Tripathi and Mishra 2021; Piczak 2015), classifying the subset of sounds that people experience in the everyday environment such as road noise or music. These background sounds can add important context to video and are robust to visual disturbances such as occlusion and illumination changes.

Audio representations have progressed significantly from the traditional signal processing representations such as FFTs or wavelets. One popular representation is mel-frequency cepstral coefficients (MFCC) which are extracted from the discrete cosine transform of the log-mel spectrum (Vimal et al. 2021). In our work, we utilize the spectrogram which is a visual representation of the frequencies across time. The spectrogram can serve as an input to convolutional neural networks or transformers. The competition, Holistic Evaluation of Audio Representations (Turian et al. 2022), conducted a head to head evaluation of the benefits of different audio representations. BYOL for Audio (Niizumi et al. 2023) tackled the problem of learning audio representations that are suitable for a multiple downstream tasks.

Since there are few labeled audio datasets, the emergence of large multimodal datasets provides fertile opportunities for leveraging cross-modality information. Several audio models have built on the success of the CLIP model at exploiting video and language. For instance, Microsoft’s CLAP (Elizalde et al. 2022) model learns a joint encoding for audio and text and can ingest text prompts to perform zero shot classification on audio data.

Prompt Learning

Performing fine-tuning on large pre-trained models like CLIP (Radford et al. 2021) leads to overfitting, which deteriorates the generalization of the model. Furthermore, it is also costly in terms of computation as all parameters of the model

need to be updated. Prompt learning is an efficient approach where the pre-trained model parameters remain frozen while introducing learnable vectors called tokens. This helps in adapting the model to downstream tasks while retaining the model’s knowledge gained during pre-training. This approach is less computationally intensive and requires less training time as compared to full fine-tuning a model. Prompt learning was introduced by researchers in the NLP area where the textual representations are similarly learned in the word embedding space (Shin et al. 2020; Jiang et al. 2020; Zhong, Friedman, and Chen 2021; Li et al. 2022). The approach was also adapted for vision tasks (Jia et al. 2022; Wang et al. 2022; Zhang, Zhou, and Liu 2022) and for usage within vision-language models (Zhou et al. 2022a; 2022b; Zhu et al. 2024). Jia et al. (2022) compare deep prompting in vision with shallow prompting and discuss the improvements in model performance due to the former contribution. Some works have discussed prompt learning in both text and vision branches (Khattak et al. 2023; Rasheed et al. 2023; Wasim et al. 2023). Khattak et al. (2023) also introduce a connection between vision and language prompts to ensure synergy while prompts are learnt.

Methodology

Our proposed multimodal architecture amalgamates the visual, text and audio transformer encoders. The former two are adapted from Vanilla CLIP (Radford et al. 2021) whereas the audio encoder is based on AudioCLIP (Guzhov et al. 2021b). During training, all encoder layers of all modalities remain frozen, therefore retaining the model’s pre-training knowledge while making it computationally efficient. To enhance the model’s capability to be effective on the downstream task of video content moderation, prompt learning is enabled for the text and vision branches. We also introduce a fully-trainable projection layer at the end of the audio branch which learns the audio representations during training. Vanilla CLIP was tailored for the video input. We calculate the aggregate representations from both audio and vision branches. The fused features along with the text features are then learned contrastively during training - maximizing the diagonal (V_i, T_i) of the cosine similarity matrix while minimizing other pairs of fused and text features. Maximizing the diagonal implies that the similarity score with the ground-truth class is maximized while others are minimized. The following sections discuss the important components of our proposed model.

Adapting CLIP for Audio and Video

Pre-training Audio Encoder The pre-trained audio head which we include in our model is ESResNe(X)t (Guzhov et al. 2021a) which is trained in four stages: 1) initialized with ImageNet weights (Deng et al. 2009), 2) fine-tuned on the AudioSet dataset (Gemmeke et al. 2017) as a standalone model, 3) trained as part of the AudioCLIP model (Guzhov et al. 2021b) keeping text and vision branches frozen, and finally 4) fully trained with all three encoders of AudioCLIP fully trainable.

in both text and vision branches of our adapted CLIP model gives the best accuracy. Similarly, increasing the number of prompt tokens in a layer improves the classification accuracy. A similar trend was observed when we increase the learning depth in terms of layers. The detailed discussion on these ablations is included in the Experiments section of this paper.

| Malicious | Benign | Total |
|-----------|--------|-------|
| 305 | 830 | 1135 |

Table 1: Distribution of malicious and benign videos in the MMOB dataset.

Experiments

The Multimodal Malicious or Benign Dataset

We curate the **Multimodal Malicious or Benign (MMOB)** dataset by selecting the samples containing malicious videos with malicious audio tracks, and benign videos with benign audio tracks. The samples have been adapted from the MOB dataset (Ahmed et al. 2023). MOB is a cartoon video dataset including visual annotations only. We extract the audio from the videos and generate the test-train splits for this new multimodal dataset. Table 1 shows the class-wise sample counts.

Training and Evaluation

We evaluate the model with the MMOB dataset in both supervised and few-shot settings. For the supervised learning setting, we use the full training split of our dataset while for the few shot learning settings, a subset of $k_i = 0, 1, 2, 4, 8, 16$ videos is randomly sampled from the training split. For all experiments, we use ViT-B/16 as our base model which is pre-trained using CLIP.

For the default setup, we employ 12 layers of prompt learning for both the CLIP Text encoder and CLIP Video encoder. Furthermore, each encoder utilizes prompt learning with 12 tokens. It’s worth noting that, in our case, the projection layer at the end of the audio head is learned during the training process.

During the training, we take 16 frames for each video along with the associated sound spectrograms and use the class label as text. We also pass on learnable prompts to both video encoder and text encoder where X represents the number of tokens used by the video encoder while Y represents the number of tokens for text encoder. For most of the experiments, the default value of 12 tokens for both video encoder and text encoder is used. However, to analyze the impact of these learnable prompts, we evaluate the effect of decreasing the number of tokens. The learning rate is set to 8×10^{-5} in all experiments. Lastly, the depth of the encoders is represented by D as depicted in Figure 3.

Results and Discussion

In this section, we present our empirical analysis using the pre-trained Vanilla CLIP model, where the base model ViT-B/16 is used, unless explicitly mentioned otherwise. Throughout our experiments, we leverage all modalities for both base models. Furthermore, we maintain the default settings for the number of tokens and the depth of text prompt layers i.e. 10 and 12, respectively. Notably, ViT-B/16 demonstrated superior performance compared to ViT-B/32, achieving an improvement of over 5.5%. The reason ViT-B/16 may exhibit better performance than ViT-B/32 can be attributed to the finer details captured by the smaller patch size. We fine-tune Vanilla CLIP with both base models for 20 epochs and report the accuracy in Table 3.

| Modalities | | | Learnable | | | Acc |
|------------|-------|-------|-----------|-------|-------|--------------|
| Text | Video | Audio | Text | Video | Audio | |
| ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | 81.49 |
| ✓ | ✓ | ✗ | ✓ | ✓ | ✗ | 78.41 |
| ✓ | ✓ | ✗ | ✓ | ✗ | ✗ | 76.65 |
| ✓ | ✓ | ✓ | ✓ | ✗ | ✓ | 78.21 |

Table 2: Adding modality and learning prompts (text/video) along with the projection layer (audio) improves the overall accuracy of the model.

Impact of Adding Audio Modality and Learnable Prompts We evaluate the improvement derived from adding the audio modality, deep learnable prompts on both vision and text branches, and enabling the learning of the projection layer deployed at the end of the audio encoder. Our approach involves systematically eliminating modalities and learnable modules, one at a time in each experiment. To better understand the impact of these configuration changes we maintain a consistent number of prompts and depth of prompt layers. The ensuing analysis of the model’s performance is presented in Table 2. As anticipated, the inclusion of additional modalities and prompt learning contributes positively to the overall accuracy. The baseline accuracy of 81.49 % in the supervised setting was observed in the configuration where we include 1) audio modality with video, 2) prompting in language and vision branches, and 3) learning of the audio projection layer.

Impact of Token Length To delve deeper into the model’s performance, we systematically diminish the number of tokens in our investigation. In this set of experiments, we maintain the utilization of all three modalities, each equipped with learnable prompt layers. The sole variation lies in the adjustment of token quantity. Table 4 provides a comprehensive summary of the results derived from these experiments. Using a greater number of tokens positively contributes to the learning of the model.

Impact of Prompt Learning Depth The role of prompt learning is pivotal in recent vision-language models. In this

| Base Model | Modalities | | | Learnable Prompts | | | No. of Tokens | | Prompt Depth | | Accuracy |
|------------|------------|-------|-------|-------------------|-------|-------|---------------|-------|--------------|-------|--------------|
| | Text | Video | Audio | Text | Video | Audio | Text | Video | Text | Video | |
| ViT-B/16 | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | 12 | 12 | 12 | 12 | 81.49 |
| ViT-B/32 | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | 12 | 12 | 12 | 12 | 77.09 |

Table 3: Results of fine-tuning CLIP on our MMOB dataset with different base models and learnable prompts. ViT-B/16 achieves better accuracy in comparison to the other model. Both models have been fine-tuned for 20 epochs.

| No. of Tokens | | Accuracy |
|---------------|-------|--------------|
| Text | Video | |
| 10 | 10 | 81.49 |
| 10 | 8 | 80.17 |
| 10 | 6 | 79.29 |
| 10 | 4 | 77.97 |

Table 4: Increasing the number of learnable prompt tokens improves performance.

series of experiments, we specifically explore the influence of prompt learning depth. Once more, we employ all three modalities, maintaining a consistent token count of 10 for both text and video heads. The outcomes of these experiments are detailed in Table 5. A discernible pattern emerges, indicating that as the depth of prompt learning increases, there is a corresponding enhancement in the model’s accuracy.

| Depth of Prompt Learning | | Accuracy |
|--------------------------|-------|--------------|
| Text | Video | |
| 12 | 12 | 81.49 |
| 8 | 8 | 80.61 |
| 4 | 4 | 77.97 |
| 2 | 2 | 75.77 |

Table 5: Multiple depth levels for prompt learning also enhance model’s performance.

Few-shot Setting Given that Vanilla CLIP is a pretrained network, it lends itself well to few-shot learning. In these experiments, we introduce the few-shot learning parameter, denoted as k , with values specified as $k_i = \{0, 1, 2, 4, 8, 16\}$. This means that for each class, we utilize k_i samples, while keeping the other parameters set to their original values. The samples are drawn uniformly from each class without replacement.

| k_i | Accuracy |
|-------|--------------|
| 0 | 59.47 |
| 1 | 63.01 |
| 2 | 64.75 |
| 4 | 63.87 |
| 8 | 61.67 |
| 16 | 54.18 |

Table 6: Few shot learning for various k samples of each class.

In few shot learning experiments, we see an interesting trend. For $k = 2$, the model achieves the highest accuracy while as the value of k increases, the accuracy drops. This trend is shown in Table 6. The choice of k affects the intrinsic dimensionality of the learned embedding space, hence increasing k does not result in monotonic performance improvements (Cao, Law, and Fidler 2019).

Conclusion and Future Work

This paper highlights the importance of leveraging the audio modality for the problem of content moderation of children’s videos. We introduce a new multimodal architecture that adds a pre-trained audio encoder with a learnable projection layer to the adapted CLIP model. Our adapted model averages video frames to create a visual representation and learns prompts on both branches of Vanilla CLIP (text and vision). This contribution enhances the overall training and performance of video content moderation. We also created a multimodal dataset, MMOB, for this task that includes audio annotations. In future work, we intend to extend our work to other types of online video content including YouTube shorts, Facebook reels, and TikTok videos. These short duration videos are popular with viewers but may create negative impacts such as reducing attention span. Lastly, we plan to work on the problem of moderating the advertisements that accompany children’s videos. From personal experience browsing these platforms, we observed that although a video may be appropriate for viewing by young children the advertisements occasionally contain unsuitable content.

References

- Ahmed, S. H.; Khan, M. J.; Qaisar, H. M. U.; and Sukthankar, G. 2023. Malicious or benign? Towards effective content moderation for children's videos. *The International FLAIRS Conference Proceedings* 36(1).
- Alghowinem, S. 2019. A safer YouTube Kids: An extra layer of content filtering using automated multimodal analysis. In Arai, K.; Kapoor, S.; and Bhatia, R., eds., *Intelligent Systems and Applications*, 294–308. Cham: Springer International Publishing.
- Cao, T.; Law, M.; and Fidler, S. 2019. A theoretical analysis of the number of shots in few-shot learning. *arXiv preprint arXiv:1909.11722*.
- Chuttur, M. Y., and Nazurally, A. 2022. A multi-modal approach to detect inappropriate cartoon video contents using deep learning networks. *Multimedia Tools and Applications* 81(12):16881–16900.
- COPPA. 1998. Children's online privacy protection act of 1998. Public Law 105-277, 112 Stat. 2681-728.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Elizalde, B.; Deshmukh, S.; Ismail, M. A.; and Wang, H. 2022. CLAP: Learning Audio Concepts From Natural Language Supervision. *arXiv:2206.04769* [cs, eess].
- FTC. 2024. The federal register.
- Gemmeke, J. F.; Ellis, D. P.; Freedman, D.; Jansen, A.; Lawrence, W.; Moore, R. C.; Plakal, M.; and Ritter, M. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 776–780. IEEE.
- Guzhov, A.; Raue, F.; Hees, J.; and Dengel, A. 2021a. Es-resne (x) t-fbsp: Learning robust time-frequency transformation of audio. In *International Joint Conference on Neural Networks (IJCNN)*, 1–8. IEEE.
- Guzhov, A.; Raue, F.; Hees, J.; and Dengel, A. 2021b. Audioclip: Extending CLIP to image, text and audio.
- Hapkiewicz, W. G. 1979. Children's reactions to cartoon violence. *Journal of Clinical Child Psychology* 8(1):30–34.
- Hayes, A. 2024. YouTube stats: Everything you need to know in 2024!
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep residual learning for image recognition.
- Hinkley, T., and McCann, J. R. 2018. Mothers' and father's perceptions of the risks and benefits of screen time and physical activity during early childhood: a qualitative study. *BMC Public Health* 18(1):1–8.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning.
- Jiang, Z.; Xu, F. F.; Araki, J.; and Neubig, G. 2020. How can we know what language models know?
- Khattak, M. U.; Rasheed, H.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. MaPLe: Multi-modal prompt learning. (arXiv:2210.03117). *arXiv:2210.03117* [cs].
- Klatte, M.; Bergström, K.; and Lachmann, T. 2013. Does noise affect learning? A short review on noise effects on cognitive performance in children. *Frontiers in Psychology* 4.
- Le, B. M.; Tandon, R.; and Oinar, C. 2022. Samba: Identifying inappropriate videos for young children on YouTube. In *Proceedings of the ACM International Conference on Information and Knowledge Management*.
- Li, Y.; Liang, F.; Zhao, L.; Cui, Y.; Ouyang, W.; Shao, J.; Yu, F.; and Yan, J. 2022. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm.
- Lillard, A. S., and Peterson, J. 2011. The immediate impact of different types of television on young children's executive function. *Pediatrics* 128(4):644–649.
- McLoughlin, I.; Zhang, H.; Xie, Z.; Song, Y.; and Xiao, W. 2015. Robust sound event classification using deep neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 23(3):540–552.
- Niizumi, D.; Takeuchi, D.; Ohishi, Y.; Harada, N.; and Kashino, K. 2023. BYOL for Audio: Exploring Pre-Trained General-Purpose Audio Representations. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 31:137–151.
- Ortutay, B. 2023. States sue Meta claiming its social platforms are addictive and harm children's mental health.
- Papadamou, K.; Papasavva, A.; Zannettou, S.; Blackburn, J.; Kourtellis, N.; Leontiadis, I.; Stringhini, G.; and Sirivianos, M. 2021. Disturbed YouTube for Kids: Characterizing and detecting inappropriate videos targeting young children. (arXiv:1901.07046). *arXiv:1901.07046* [cs].
- Piczak, K. J. 2015. ESC: Dataset for Environmental Sound Classification. In *Proceedings of the 23rd Annual ACM Conference on Multimedia*, 1015–1018. ACM Press.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning transferable visual models from natural language supervision.
- Rasheed, H.; Khattak, M. U.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. ViFi-CLIP: Fine-tuned CLIP models are efficient video learners. (arXiv:2212.03640). *arXiv:2212.03640* [cs].
- Shewale, R. 2024. Video marketing statistics in 2024 (usage, roi and more).
- Shin, T.; Razeghi, Y.; Logan IV, R. L.; Wallace, E.; and Singh, S. 2020. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In Webber, B.; Cohn, T.; He, Y.; and Liu, Y., eds., *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4222–4235. Online: Association for Computational Linguistics.
- Tahir, R.; Ahmed, F.; Saeed, H.; Ali, S.; Zaffar, F.; and Wilson, C. 2020. Bringing the kid back into YouTube Kids: detecting inappropriate content on video streaming platforms. *ASONAM '19*, 464–469. New York, NY, USA: Association for Computing Machinery.

TikTok. 2024. Our approach to content moderation.

Tripathi, A. M., and Mishra, A. 2021. Environment sound classification using an attention-based residual neural network. *Neurocomput.* 460(C):409–423.

Turian, J.; Shier, J.; Khan, H. R.; Raj, B.; Schuller, B. W.; Steinmetz, C. J.; Malloy, C.; Tzanetakis, G.; Velarde, G.; McNally, K.; Henry, M.; Pinto, N.; Noufi, C.; Clough, C.; Herremans, D.; Fonseca, E.; Engel, J.; Salamon, J.; Esling, P.; Manocha, P.; Watanabe, S.; Jin, Z.; and Bisk, Y. 2022. HEAR: Holistic Evaluation of Audio Representations.

Vimal, B.; Surya, M.; Darshan; Sridhar, V.; and Ashok, A. 2021. MFCC based audio classification using machine learning. In *International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 1–4.

Wang, Z.; Zhang, Z.; Ebrahimi, S.; Sun, R.; Zhang, H.; Lee, C.-Y.; Ren, X.; Su, G.; Perot, V.; Dy, J.; and Pfister, T. 2022. Dualprompt: Complementary prompting for rehearsal-free continual learning.

Wasim, S. T.; Naseer, M.; Khan, S.; Khan, F. S.; and Shah, M. 2023. Vita-CLIP: Video and text adaptive CLIP via multimodal prompting. (arXiv:2304.03307). arXiv:2304.03307 [cs, eess].

YouTube. 2024. How does YouTube manage harmful content?

Zhang, Y.; Zhou, K.; and Liu, Z. 2022. Neural prompt search.

Zhong, Z.; Friedman, D.; and Chen, D. 2021. Factual probing is [mask]: Learning vs. learning to recall.

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Co-coop: Conditional prompt learning for vision-language models. (arXiv:2203.05557).

Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Coop: Learning to prompt for vision-language models. *International Journal of Computer Vision* 130(9):2337–2348.

Zhu, B.; Niu, Y.; Han, Y.; Wu, Y.; and Zhang, H. 2024. Prompt-aligned gradient for prompt tuning.