

# Automatic Funny Scene Extraction from Long-form Cinematic Videos

Sibendu Paul, Haotian Jiang, Caren Chen

Amazon Prime Video, Seattle, USA  
{sibendu, haotij, carechen}@amazon.com

## Abstract

Automatically extracting engaging and high-quality humorous scenes from cinematic titles is pivotal for creating captivating video previews and snackable content, boosting user engagement on streaming platforms. Long-form cinematic titles, with their extended duration and complex narratives, challenge scene localization, while humor’s reliance on diverse modalities and its nuanced style add further complexity. This paper introduces an end-to-end system for automatically identifying and ranking humorous scenes from long-form cinematic titles, featuring shot detection, multimodal scene localization, and humor tagging optimized for cinematic content. Key innovations include a novel scene segmentation approach combining visual and textual cues, improved shot representations via guided triplet mining, and a multimodal humor tagging framework leveraging both audio and text. Our system achieves an 18.3% AP improvement over state-of-the-art scene detection on the OVSD dataset and an F1 score of 0.834 for detecting humor in long text. Extensive evaluations across five cinematic titles demonstrate 87% of clips extracted by our pipeline are intended to be funny, while 98% of scenes are accurately localized. With successful generalization to trailers, these results showcase the pipeline’s potential to enhance content creation workflows, improve user engagement, and streamline snackable content generation for diverse cinematic media formats.

## Introduction

With the rapid growth of video-streaming platforms, such as Netflix and Prime Video, there is an increasing need to enhance user engagement through personalized and engaging content. With Netflix boasting 300 million subscribers (15.9% YoY growth) and Prime Video over 200 million (10% annual growth), these platforms increasingly rely on personalized video ads, and short-form video content to enhance user experience. One such avenue is the creation of short, humorous video clips that can capture user attention and improve discoverability. Humor, an essential element of human interaction, can be a powerful tool in promotional content, capturing attention (Garner 2005; Wanzer, Frymier, and Irwin 2010), building rapport (Stauffer 1999), and fostering trust (Vartabedian and Vartabedian 1993). However, the current process of extracting high-quality humorous scenes from

full-length cinematic titles is predominantly manual, requiring significant time and effort. To address this, we propose an end-to-end automated pipeline for funny scene extraction that leverages multimodal data and deep learning to reduce costs and expand coverage across movies, web series, and trailers. The pipeline analyzes shots<sup>1</sup> from any cinematic title, groups them into semantically cohesive scenes, and then identifies humorous moments for efficient extraction of engaging, contextually relevant content — enhancing promotional content creation and supporting the marketing goals of streaming platforms.

We address three key challenges in extracting humorous scenes from long-form cinematic content: (1) robust scene detection over complex temporal structures, (2) accurate humor identification within extended scenes, and (3) filtering of improper humor with reliable ranking of genuine humorous segments. Unlike short-range video analysis tasks such as shot detection or action recognition, scene localization in long-form videos requires reasoning over both short- and long-range dependencies, yet suffers from scarce large-scale annotations due to the high cost of labeling boundaries. While self-supervised learning has advanced for images and short clips, its extension to long-form content is hindered by simplistic data curation and limited multi-modal signal extraction.

To address these challenges, we propose a multi-modal scene-segmentation framework that integrates visual and text signals, operating even without audio or with only sparse key frames per each shot. For visual signal, we draw inspiration from face-clustering (Schroff, Kalenichenko, and Philbin 2015) to group similar shots into scenes, using triplet loss to capture contextual relationships. Unlike prior contrastive shot-representation methods (Qian et al. 2021; Chen et al. 2021; Mun et al. 2022) that rely on artificial augmentations or heuristic pair generation and ignore scene structure, we introduce a *guided triplet generation* strategy leveraging MovieNet-SSeg (Huang et al. 2020) to improve intra-scene similarity and inter-scene discrimination. To address its sparse key-frame coverage, we incorporate shot-level captions as complementary text cues, enabling robust representation even without audio. This multi-modal approach

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>A shot is a sequence of frames captured continuously by the same camera (Thirard and Codelli 1994)

generalizes effectively to benchmarks such as OVSD (Rotman, Porat, and Ashour 2016), demonstrating its suitability for real-world long-form scene localization.

Aside from the scene localization, humor detection in long-form cinematic content is challenging due to its multi-faceted nature, spanning words, gestures, prosodic cues, and contextual dependencies. Prior work largely focuses on single-modal approaches for short texts or videos (Yang et al. 2015; Chen and Lee 2017), limiting applicability to extended scenes. We propose a multi-modal humor-tagging pipeline that fuses audio cues (prosody, laughter) with textual understanding of context–punchline relationships, using a modified ColBERT architecture with tailored training and deployment for long text. To safeguard user experience, we integrate a guardrail audio-tagging model to filter improper humor such as bullying or mockery. Finally, we introduce a heuristic-based humor scoring mechanism to rank scenes, enabling streaming platforms to surface the most engaging promotional content.

Overall, our contributions can be summarized as follows:

- We introduce an end-to-end automatic pipeline for extracting funny scenes from long-form cinematic titles.
- We propose a multimodal scene segmentation method that encodes visual and textual cues from individual shots, leveraging triplet loss with MovieNet-SSeg supervision to improve shot–scene associations.
- We design a novel shot encoder combining the X-CLIP cross-frame transformer with a DINO projection head, enabling efficient representation learning with minimal training (80K triplets over 25 epochs).
- Our scene detection model achieves state-of-the-art performance on MovieNet-SSeg and improves AP on OVSD by 18.3%, demonstrating strong generalization.
- For humor identification, we develop a multi-modal audio-text model that achieves an F1 score of 0.834 and accuracy of 0.728, and integrate a guardrail audio-tagging model with 100% recall to filter improper humor.
- We introduce a heuristic-based humor scoring mechanism to rank scenes, with curator evaluations confirming 87% humor detection accuracy and 98% scene localization in full-length titles.

The primary use case of this end-to-end pipeline is to extract humorous scenes for autoplay functionality, enabling customers to preview content with a funny clip when hovering over a title, enhancing engagement. In the future, these output clips could be leveraged for a “fast laugh” experience on mobile platforms. Versatile by design, the pipeline can adapt to various scene extraction needs, catering to diverse streaming audience preferences.

## Related Works

**Scene Detection.** Prior Unsupervised methods (Rui, Huang, and Mehrotra 1998) cluster shots using visual features, while supervised approaches (Rao et al. 2020) improve performance but rely on costly annotations and multiple pre-defined shot features. Self-supervised methods (Dorkenwald et al. 2022; Chen et al. 2021; Mun et al. 2022) learn shot representations without annotations but still require scene boundary labels for fine-tuning alongside carefully crafted pretext

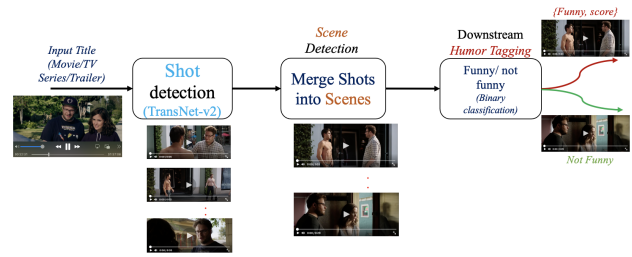


Figure 1: *Funny scene extraction pipeline overview.* The pipeline consists of three main blocks: shot detection, shot representation extraction, and scene merging, followed by downstream binary tagging and humor scoring.

tasks during pretraining. We use available scene boundaries to mine triplets and apply triplet loss for shot contrastive pre-training.

**Humor detection.** Earlier NLP research focused on humor prediction from text alone (Annamoradnejad and Zoghi 2024; Purandare and Litman 2006; Kiddon and Brun 2011), while recent work (Kayatani et al. 2021; Hasan et al. 2019; Brown, Kalogeiton, and Zisserman 2021) highlights the value of multimodal approaches. Although datasets like UR-FUNNY (Hasan et al. 2019), MHD (Patro et al. 2021), and MusTARD (Castro et al. 2019) enable such studies, they lack humor annotations for longer scenes.

**Video retrieval.** With the growth of streaming platforms, video retrieval has become vital for browsing and searching large archives. It includes video–text retrieval—matching videos to captions—explored by works like ClipBERT (Lei et al. 2021), Frozen (Bain et al. 2021), CLIP4Clip (Luo et al. 2022), DRL (Wang et al. 2024), and JSFusion (Yu, Kim, and Kim 2018), and similarity-based retrieval, such as copy detection (Jiang, Jiang, and Wang 2014) and near-duplicate retrieval (Song et al. 2011; Wu, Hauptmann, and Ngo 2007). Our focus is on extracting video scenes that preserve contextual meaning, particularly humor events — a task demanding nuanced understanding and intricate contextual reasoning, far more challenging than retrieving clips by simple tags like celebrity, place, or action.

## Methodology

Our proposed pipeline (shown in figure 1) processes various titles e.g., movies, tv episodes etc, extracting funny scenes of variable length with assigned scores. Comprising three primary blocks, our pipeline first extracts shots, then merges semantically similar shots into scenes, and finally conducts binary humor classification and assign scores.

To extract frames captured from the same camera (i.e., shots), we utilize the state-of-the-art pretrained TransNetV2 network (Souček and Lokoč 2020) as our shot-detector.

### Scene Detection

A scene is defined as a sequence of shots depicting a semantically cohesive story segment. Unlike shot detection,

which relies on visual cues, scene segmentation is more challenging due to the complex temporal relationships between shots (Chen et al. 2021). We propose a multimodal scene segmentation method that extracts and fuses visual and textual cues from each shot using separate encoders. The visual encoder is trained with self-supervised methods on task-specific data and pre-trained with Triplet Loss, while BLIP-2 (Li et al. 2023) serves as the textual feature extractor. Combined shot features are refined through supervised fine-tuning with neighboring context to accurately detect scene boundaries.

**Shot Contrastive pre-training** We employ contrastive learning to obtain shot representations that capture local scene structures for scene boundary detection. Unlike prior methods, we treat shot-level representation learning as a multi-class classification pretext task, where each scene is a class and its shots are samples per class. Using Triplet Loss — a proven success in image-based face recognition (Schroff, Kalenichenko, and Philbin 2015), we learn distinctive shot embeddings suited for scenarios where the number of classes ( $N_{Class}$ ) significantly outnumbers the number of samples for each class ( $N_{SamplesPerClass}$ ). This parallels the face-person relationship in images with the shot-scene association in videos.

**Triplet Mining.** Triplet Loss, though less greedy than other contrastive losses, depends heavily on selecting suitable anchor ( $s_i^a$ ), positive ( $s_i^p$ ) and negative ( $s_i^n$ ) (Sarigöz 2022) samples. To generate triplets, we adopt an offline approach, leveraging ground-truth scene boundaries from MovieNet-SSeg train set (Huang et al. 2020): shots within the same scene form positives, shots from different scenes form negatives, completing the triplet. To generate the hard triplets, we sample negative shots from the neighboring scenes within a  $\pm 3$ -scene window. From 190 movies, we randomly sample 420 triplets per movie, yielding  $\sim 80K$  triplets to pre-train our feature extractor. During pre-training, our goal is to make a shot ( $s_i^a$ ) from a specific scene is closer to all other shots ( $s_i^p$ ) from that scene than to any shot ( $s_i^n$ ) from any other scene in the feature space, as in Equation 1, where  $\alpha$  denotes the margin enforced between positive and negative pairs. We use a fixed  $\alpha$  value of 1 throughout our training.

$$\|f(s_i^a) - f(s_i^p)\|_2^2 + \alpha < \|f(s_i^a) - f(s_i^n)\|_2^2, \quad \forall (s_i^a, s_i^p, s_i^n) \in T \quad (1)$$

We generate triplets from ground-truth labels for added supervision, rather than using pseudo labels, leaving online hard triplet mining and progressive sampling for future work.

**Shot Visual Encoder network.** Our shot encoder learns shot-level representations from visual modalities using the XCLIP visual encoder (Ni et al. 2022) that leverages cross-frame attention to capture long-range temporal dependencies. To refine the transformer-extracted video feature representation, we append the DINO projection head (Caron et al. 2021) — a 3-layer MLP with GELU activations, L2-normalization, and a weight-normalized fully connected layer with 4096 dimensions. Hence, our shot encoder combines XCLIP’s cross-frame communication transformer with the DINO projection head, producing a 4096-dimensional feature per shot

(Figure 2).

**Data Augmentation.** In video processing, shots require both temporal and spatial augmentations. For temporal augmentation, we use sparse sampling (Wang et al. 2016) to reduce redundancy. For spatial augmentation, we follow (Dave et al. 2022) with color jitter (contrast, hue, brightness, saturation), horizontal flipping, and spatial masking inspired by VideoMAE (Tong et al. 2022) to enhance shot encoder training. The impact of these augmentations on contrastive pre-training is examined in the *Experiments* section.

**Exploiting Text Modality.** For the large-scale scene detection dataset, where only a few key frames per shot are available, to explore the text-modality during the training, we use captions generated by vision-language model, video-llava (Lin et al. 2023) as the text input for each shot. These are encoded using a pre-trained BLIP-2 (Li et al. 2023) text extractor with mean pooling, yielding 768-dimensional text features. Each shot is then represented by concatenating the 4096-dimensional visual features with the 768-dimensional text features, forming a 4864-dimensional embedding.

**Supervised finetuning** In long-range video analysis, determining scene boundaries requires considering neighboring shots. We use a window (W) covering N preceding and N succeeding shots, along with the center shot. Visual features from a triplet-loss trained shot encoder and textual features from BLIP-2 are concatenated per shot, then flattened into a  $(2 \times N + 1) \times 4864$  vector, fed to four-layer MLP SBD head ( $[2 \times N + 1] \times 4864 - 8192 - 4096 - 1024 - 2$ ), with dropout applied after each of the first three FC layers. The window slides shot-by-shot, marking the center shot as a boundary if its score exceeds 0.5. During finetuning, the shot encoder is frozen and only the MLP is trained with cross-entropy loss.

### Downstream Scene Tagging

Next, we label each scene to determine its humor content, focusing on eliciting positive emotions like amusement and laughter (Atlantic 2018). Our humor tagging approach utilizes audio and text analysis, detecting laughter in audio and using AWS-Transcribe with a custom-trained colBERT model to identify humor from verbal content. Audio from selected scenes after both analysis is further checked by an audio-tagging model to exclude scenes with improper humor. The filtered results are merged, and each humorous scene is scored to produce a ranked list of top humorous moments, as shown in Figure 3.

Recognizing the significance of laughter in identifying humor, we extract audio and detect it using a ResNet18-based model trained on the Switchboard dataset (Godfrey and Holliman 1997). We finetuned the model to accurately detect laughter from both individuals and background audiences.

**Humor detection from verbal modality.** According to Cambridge Dictionary, humor entails “the ability to be amused by something seen, heard, or thought about, sometimes causing you to smile or laugh, or the quality in something that causes such amusement”. Not all instances of amusement result in laughter, yet they may still be perceived as funny. This insight motivates us to extract humor from the

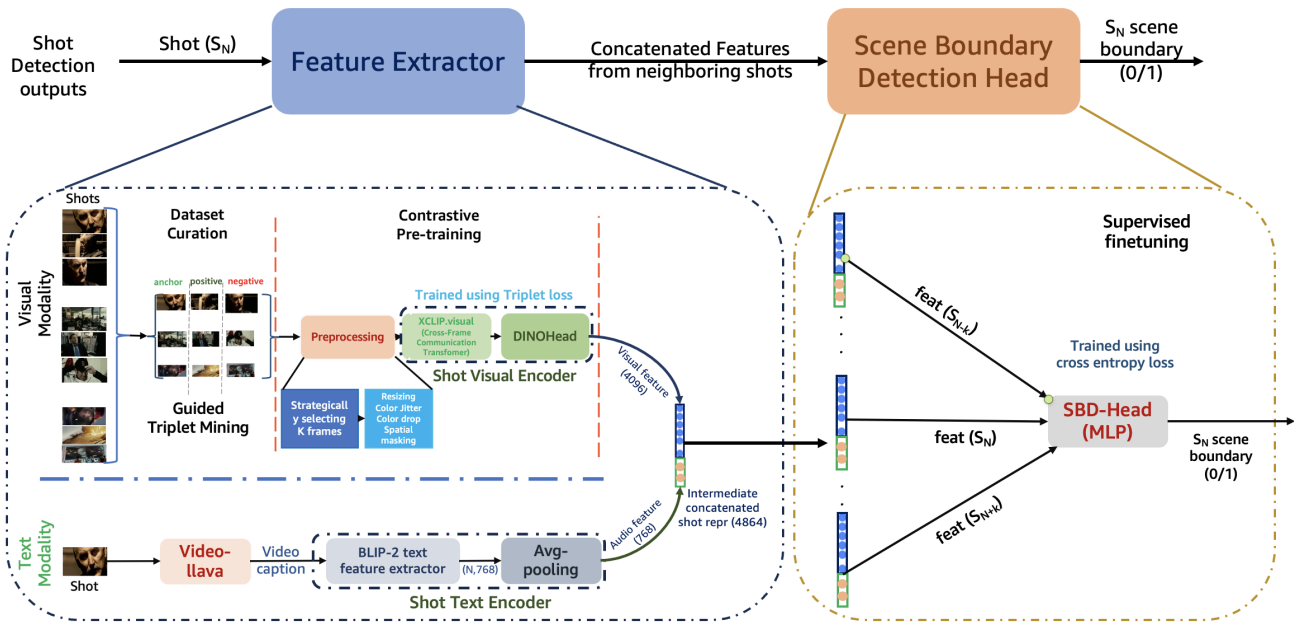


Figure 2: *Scene detection module overview*: Triplets generated from MovieNet-SSeg ground-truth boundaries; spatial and temporal augmentations applied during contrastive pretraining with triplet loss; for each video, we extract shot-level video caption and use shot text-encoder to extract shot-level text features; fine-tuning aggregates neighboring shot features and trains MLP layers via supervised learning.




Figure 3: *Multi-modal humor tagging pipeline overview*. (1) Extract audio and detect laughter, (2) transcribe audio and analyze funny conversations, (3) filter improper humor, then score and output the ranked funny scene videos.

verbal or text modality. Linguistically, humor often follows a setup, multi-stage storytelling, and a punchline that subverts expectations (Eysenck 1942; Suls 1972). Following this, the ColBERT model (Annamoradnejad and Zoghi 2024) encodes individual sentences and the full text, concatenates them, and processes them through MLP layers to capture relationships among sentences and also scrutinize word-level connections throughout the entire text to assess the global congruity.

**Model architecture and training.** The state-of-the-art ColBERT processes the first 5 input sentences along with the entire text, which works for short texts but is less effective for longer scenes (30s–2min) from cinematic titles. To address this, we adapt ColBERT by training on the UR-FUNNY dataset (Hasan et al. 2019) which has the longest average scene duration (33.8s) among humor datasets. Our version processes 10 sentences plus the full text to capture broader context–punchline relationships, using effective sampling to improve inter-sentence relationship learning.

**Training Scenario.** Inspired by the inherent structure of humor, which typically consists of a setup and punchline, we developed a 10-sentence sampling approach for training. Our semi-deterministic method handles longer texts by preserving the first two sentences for context and the last two for the punchline, while randomly selecting six sentences from the middle with appropriate spacing. This structure-dependent sampling strategy demonstrated superior generalization performance compared to other tested approaches.

**Testing Scenario.** During the testing phase, we maintain determinism by avoiding random sampling and instead segment longer texts into multiple subtexts of 10 sentences each. These subtexts are processed individually through our custom-trained ColBERT model (trained on UR-FUNNY dataset), and their softmax scores are averaged to produce the final classification score. We employ an empirically determined softmax threshold of 0.56 for humor detection to minimize false positives.

**Filtering out improper humor.** Merely detecting laughter or humorous context alone is insufficient to ascertain the humorous quality of a scene, as it might include instances of improper humor such as bullying, mocking, or taunting. To tackle this, we employ a state-of-the-art pretrained audio-tagging CNN model (Kong et al. 2020) trained on the AudioSet dataset. Using this model, we filter out inappropriate content by detecting scenes containing distress signals (crying, screaming) or negative emotional cues, thereby ensuring only appropriate humorous content reaches viewers<sup>2</sup>.

<sup>2</sup>Our current guardrail focuses on audio-based moderation, which effectively captures inappropriate humor violations such as

Shot-Encoder Variants	V1	V2	V3	final
<i>NMI</i> -overall	0.606	0.604	0.611	<b>0.632</b>

Table 1: Scene clustering quality evaluated using normalized mutual information (NMI) across various triplet mining strategies.

**Humor Score Generation.** We developed a heuristic scoring system for funny scenes using curator feedback as ground truth. It combines four normalized features—average laughter ( $f_1$ ), laughter duration above threshold ( $f_2$ ), ColBERT softmax score ( $f_3$ ), and scene length ( $f_4$ )—weighted via grid search (Equation 2). This enables effective ranking of top-k scenes against curator annotations for prioritizing humorous content for streaming media customers.

$$humor_{score} = (w_1 \times f_1) + (w_2 \times f_2) + (w_3 \times f_3) + (w_4 \times e^{-f_4/t_c}) \quad (2)$$

We acknowledge that humor is subjective and culturally variable, and our heuristic scoring may not fully capture styles like sarcasm, wordplay, or region-specific comedic cues. In future work, we aim to develop learned ranking models trained on more diverse, demographically varied annotations to improve cross-cultural generalization.

## Experiments

### Evaluation Metrics

We evaluate each subtask independently using standard metrics: F1 for shot detection, AP and F1 for scene localization, and binary accuracy with F1 for our custom-trained colBERT model. Due to limited end-to-end evaluation data, we conduct a user study with professional curators, reporting both binary accuracy and qualitative insights.

### Experimental Results

Our pipeline comprises three sequential blocks (Figure 1), with the optimal model selected for each stage and earlier stages fixed to their best-performing versions. Experimental results report the performance of each stage individually.

We observed that our shot detector, TransNetV2 (Souček and Lokoč 2020), outperforms AUTOSHOT (Zhu et al. 2023) by 1.2% on average F1 metric, even after a finetuned threshold on the BBC Planet Earth video dataset (Baraldi, Grana, and Cucchiara 2015).

**Scene Detection** We here use Average Precision (AP)<sup>3</sup> and F1 score (Chen et al. 2021).

**Impact of guidance while triplet generation.** To evaluate the effect of guidance in triplet mining on shot-encoder training, we compare three random heuristic-based strategies that disregard annotations. In V1, positive shots are chosen

violent cues. However, we acknowledge that a comprehensive safety system would benefit from multimodal guardrails that incorporate visual and textual modalities.

<sup>3</sup>AP value condenses the precision-recall (PR) curve into a single, confidence-threshold-independent metric that reflects the average of all precision values, making it well-suited to our use case.

Version No.	Augmentations					best AP	best F1
	Resize	Color Jitter	Horizontal Flip	Random Masking	Uniform Masking		
A1	✓	✓	-	-	-	41.7	35.1
<b>A2</b>	✓	✓	✓	✓	-	<b>47.2</b>	<b>41.8</b>
A3	✓	✓	✓	-	8%-15%	38.1	16.8
A4	✓	✓	✓	-	16%-25%	34.3	15.2

Table 2: Impact of different augmentations on Movienet-SSeg.

Model	best AP
ShotCoL (Chen et al. 2021)	25.02
BASSL (Mun et al. 2022)	28.68
<b>Ours (Visual)</b>	<b>33.94</b>

Table 3: Comparison between our method and other shot-level pre-training baselines on OVSD, utilizing visual signals (V) only.

within a  $\pm 3$ -shot window and negatives are  $\geq 10$  shots away; in V2, the window is reduced to  $\pm 2$  shots and negatives are  $\geq 15$  shots away; in V3, inspired by MovieNet-SSeg statistics, positives are within  $\pm 1$  shot and negatives are  $\geq 30$  shots away. Each variant is used to pre-train the shot encoder, and clustering quality is measured using normalized mutual information (NMI) over 100 randomly sampled scenes from the MovieNet-test split. As shown in Table 1, V3 yields fewer false positives and better clustering than V1 and V2, while our guided triplet generation - eliminating false-positive triplets and achieves the highest NMI, highlighting the value of incorporating guidance in triplet mining for improved shot-representation learning.

**Impact of different augmentations.** We apply various data augmentations during shot contrastive pretraining to enhance representation robustness. Table 2 compares four combinations (A1–A4), with A2—featuring resizing, color jitter, horizontal flipping, and random masking—achieving the best scene detection accuracy. The masking strategy, inspired by prior work, improves recall but reduces precision, with A2 striking the best balance to boost AP and F1.

**Generalization performance.** MovieNet-SSeg (Huang et al. 2020) provides 42K scenes from 318 movies with 3 key frames per shot, while OVSD (Rotman, Porat, and Ashour 2016) offers continuous video input, aligning closely with our (Prime Video) deployment scenario. We extract 8 uniformly spaced frames per shot to reduce temporal redundancy. Table 3 shows that our model, trained on MovieNet-SSeg using only visual features, generalizes well to OVSD without finetuning.

**Impact of incorporating Text Modality.** Incorporating text-based features (768) alongside larger visual features (4096) significantly improves scene detection. As shown in Table 4, adding text boosts AP by 9.1% and F1 by 11% on

the MovieNet-SSeg test set, enabling the model to leverage unique shot-level cues for better scene localization.

Input-modality	best AP	best F1
Visual (V)	47.2	41.8
Visual (V) + Text (T)	51.5	46.4

Table 4: Impact of input modalities on the scene detection performance on Movienet-Sseg test set.

**Downstream Humor Tagging** After optimizing the first two stages of the pipeline depicted in Figure 1, we evaluate the performance of the custom-trained colBERT model for humor tagging based on text modality. Due to space constraints, the individual impact of the audio modality is omitted here.

**Using Text Modality.** We evaluated four model variants based on input sentence count (5 or 10) and embeddings (bert-base-uncased or roberta-base). The colBERT model trained on the UR-FUNNY dataset (Hasan et al. 2019) using 10 input sentences and bert-base embeddings achieved the best generalization performance on the MHD dataset (Patro et al. 2021), with an F1 score of  $0.834$  and an accuracy of  $0.728$ . This outperforms other variants, the state-of-the-art finetuned transformer (Weller and Seppi 2019), and even multimodal fusion models like FunnyNet (Liu, Courant, and Kalogeiton 2022), despite using only text modality.

**Evaluation of Humor Score Generation.** To rate the humor in output scenes for viewer presentation, we propose a heuristic-based mechanism relying on four high-level signals: (1) average laughter score ( $f_1$ ), (2) percentage of time laughter score exceeds a threshold ( $f_2$ ), (3) custom-trained ColBERT model softmax score ( $f_3$ ), and (4) scene duration ( $f_4$ )<sup>4</sup>. The weights  $[w_1, w_2, w_3, w_4]$  (used in equation 2) are optimized using grid search across three regression models: linear, logistic, and decision-tree regression.

For evaluation, we propose a custom intersection-over-union (IoU)-based metric combining coarse-grained and fine-grained estimations. The coarse-grained metric,  $topIOU_N$ , measures overlap between the top-N humorous scenes predicted by our pipeline and those identified by curators<sup>5</sup> (Equation 4). The fine-grained metric,  $topIOU\_Align_N$ , evaluates alignment between the top-N predicted funny scenes and the curators’ top-N ranked scenes (Equation 5). Both metrics are normalized for three N values (3, 5, 10) and aggregated into an overall evaluation metric,  $eval_{metric}$ , as shown in Equation 3.

$$eval_{metric} = \sum_{N=\{3,5,10\}} (topIOU_N + topIOU\_Align_N) \quad (3)$$

$$topIOU_N = |\{gt[:]\} \cap \{predicted[: N]\}|/N \quad (4)$$

$$topIOU\_Align_N = |\{gt[: N]\} \cap \{predicted[: N]\}|/N \quad (5)$$

<sup>4</sup>Shorter funny scenes are preferred by professional curators.

<sup>5</sup>Output funny scenes scored and ranked by professional curators serve as the ground truth i.e.,  $\{gt\}$ .

Using 60% of curator-annotated samples for weight estimation and testing on the remaining 40%, we found logistic regression achieves the best  $eval_{metric}$  score on the test set.

Train Set	Test Set	Model	Acc	F1
UR-FUNNY	MHD	Fine-tuned Transformer	0.591	0.719
		ColBERT-base-5	0.694	0.820
		<b>ColBERT-base-10</b>	<b>0.728</b>	<b>0.834</b>
		ColRoBERTa-base-5	0.720	0.830
		ColRoBERTa-base-10	0.690	0.812
MHD	MHD	MASM (Patro et al. 2021)	0.720	0.810

Table 5: Our text-based humor tagging model performance.

### Professional curators’ Evaluation

To comprehensively evaluate our funny scene generation pipeline, we tested it on five  $\sim 2$ -hour films spanning comedy, drama, action, adventure, and animation, as well as on 11 trailers from diverse genres. The pipeline (Figure 1) produced  $\sim 100$  scenes from the films and 17 scenes (10–15s each) from the trailers, all reviewed by three professional curators. Curators assessed each clip for objective humor (clear audio-visual or textual cues such as in-context laughter, smiles, or on-screen jokes, excluding improper content like bullying or cultural stereotypes, and including subtle forms like sarcasm and witty remarks), subjective humor (whether the clip felt funny regardless of explicit cues or creator intent), and scene ending quality (a natural, complete conclusion without abrupt cuts or unresolved dialogue).

Results (Table 6) show 98% scene localization accuracy for main content, with higher segmentation errors in trailers due to rapid transitions. Under objective guidelines, humor detection accuracy reached 87% for films and 100% for trailers; subjective evaluation yielded 74% and 88%, respectively. While trailers had more segmentation errors, their condensed, event-rich nature made humor easier to detect. Inter-annotator agreement, measured by pairwise Cohen’s  $\kappa$ , had a minimum value of 0.75, indicating substantial agreement among curators in assessing the humorous nature of the clips. Our approach extracts humorous scenes independent of genre, and we currently use a simple heuristic that skips the final 20% of a title to avoid spoilers. A more sophisticated spoiler-aware mechanism would further improve robustness across all genres.

	Movies	Trailers
Is the clip intended to be funny?	87%	100%
Do curators find the clip funny?	74%	88%
Proper Scene Ending?	98%	53%

Table 6: curators’ evaluation on movies & trailers.

**Qualitative Analysis.** Here, we present example funny scenes extracted by our proposed pipeline using multimodal




**Transcript:** Your security badges, please, this way. *Hey, listen, you're my biggest fan. Good to see you. I'm, I'm your biggest.*

(a) Output funny scene based on funny conversation detection.



(b) Laughter and funny conversation detected with “moaning, grunting” audio labels (instance of bullying).



**Transcript:** OK, I'm just the person to fix it. Yeah. You watch that I'll watch this Oh, you must be dying to know about the science fair. The results won't be posted till tomorrow, but I gave a spectacular presentation. Well, maybe I'll be a science major if I get into Felder. You will, Alex. That is, if you can pull off the 3.0, I promise them this semester. Hayley, look right there. OK. Ready? I wish I really wish I were Alexia Wheaton. Yes, I'm home. How do you feel? How tall. Oh, how am I how am I ever gonna keep this meat? You'll get used to it. I'll help you. Come on.

(c) Output funny scene based on laughter detection.



**Transcript:** Parents own the happy platter. Kid works there part-time. Happy platter? Why do we drive all the way across town to the happy platter. We'd like a booth over there near the Philodendron. Good, right? Near the Philodendron. *And this platter doesn't look all that happy to me. It looks bored, the bored platter.* I thought I would want to change the pace from drive-in food. I like drive-in food. Does this mean vegetables? A balanced diet means vegetables, kiddo. Get used to it. Good evening everyone. Oh, all right? Normally she doesn't ever drip like this. Would you like water, sir? Yes, yes, I would. This is my daughter who you must know, right? Hello Violet. Hello Violet. Hey, bye, say hi to him. Don't push it, Dad. I'm, her little brother. Hello, that's really good water. It's very refreshing spring water, is it? I don't know, sir. I think it's tap. Well, it is very good. Acts like tap. Excuse me. Well, nice to meet you. Where did she go? to find a good place to be angry.

(d) Output funny scene based on both laughter and funny conversation detection.

Figure 4: Example extracted humorous scenes and instances of improper humor from various long-form cinematic titles. Observed laughter is marked with orange boxes, while laugh-worthy dialogue or humorous context is highlighted in violet.

cues. Figure 4a and 4d show funny scenes from an action/adventure anime movie, detected using our multimodal pipeline. In Figure 4a, humor arises from incongruity as a character contradicts himself—first calling a guest his fan, then claiming to be the guest’s biggest fan. In Figure 4d, the pipeline captures both laughter and witty dialogue, including mocking a restaurant named “Happy Platter” as “Bored Platter” and joking about its water. Figure 4c presents a comedy/family movie scene identified through repeated laughter cues. In contrast, Figure 4b highlights improper humor, where students mock another’s distress, detected via audio cues.

## Pathway to Deployment & Future Work

The end-to-end automatic humor scene extraction pipeline is being deployed in phases, with TransNet-v2 shot detection and multi-modal scene detection modules already operational and supporting several use cases, including the generation of compelling, socially shareable personalized scenes designed to boost Prime Video user engagement. The final phase involves deploying the humor tagging pipeline, which requires meticulous testing and robust guardrails to ensure high precision in identifying appropriate humorous content, prioritizing accuracy over recall to maintain a positive viewing experience for Prime Video customers.

Our method performs strongly on long-form content but struggles with rapid transitions in trailers; future work will explore adaptive window sizes and enhanced temporal modeling to address this. We will also investigate progressive triplet-hardness scheduling to improve shot-representation

generalization.

While our ColBERT humor detection model currently supports only English language content, we plan to expand into multilingual humor detection and conduct comprehensive user studies to minimize evaluation bias and enhance global accessibility.

## Discussion & Future work

Our proposed multimodal scene segmentation performs well for long-form cinematic titles but faces challenges in fast-paced trailers due to rapid transitions. To address this, we propose adaptive window sizes for neighboring shot representations or enhanced temporal modeling. The adaptive approach employs a dynamic moving window (W) with multi-scale sizes (1, 2, and 4), prioritizing immediate neighbors while accommodating varying scene lengths. This method improves boundary detection across diverse shot dynamics.

Enhanced temporal modeling can be achieved through improved dataset curation during training. Our visual shot contrastive pre-training uses guided triplet generation with varying hardness levels. Selecting the hardest negatives early can cause model collapse due to bad local minima. To counter this, we propose gradually increasing triplet hardness by reducing the  $\alpha$  value in equation 1 during training (Schroff, Kalenichenko, and Philbin 2015). While this adds complexity, online triplet mining enhances scene representation learning by improving tolerance to intra-class variance and preventing embedding space over-clustering (Xuan, Stylianou, and Pless 2020; Sikaroudi et al. 2020), leading to better generalization. Future work will explore these strategies to improve scene segmentation and enhance funny scene extraction performance.

Due to the lack of annotated datasets for end-to-end evaluation, we rely on assessments by professional curators. As humor is inherently subjective, this may introduce bias; however, we mitigate this through standardized evaluation protocols and consistent guidelines. Future work will include broader user studies to improve reliability.

While laughter detection and audio-based improper humor filtering support non-English titles, our text-based funny conversation detection currently relies on English transcriptions. Since humor varies across languages, future work will focus on designing language-specific humor detection models to support global streaming catalogs.

## Conclusion

In this work, we present an end-to-end system for automatically extracting humorous scenes from long-form cinematic content to generate engaging previews. By combining shot detection, multimodal scene localization using visual and caption-based representations, and multimodal humor tagging, the system significantly improves performance – achieving an 18.3% gain on the OVSD dataset and surpassing state-of-the-art humor detection baselines. Professional curator evaluations further validate the approach with 87% accuracy on main content, while the modular design enables easy adaptation to diverse scene extraction and personalization tasks.

## References

- Annamoradnejad, I.; and Zoghi, G. 2024. ColBERT: Using BERT sentence embedding in parallel neural networks for computational humor. *Expert Systems with Applications*, 123685.
- Atlantic, T. 2018. What Makes Something Funny? <https://www.psychologicalscience.org/news/what-makes-something-funny-2.html>. Accessed: 2025-12-30.
- Bain, M.; Nagrani, A.; Varol, G.; and Zisserman, A. 2021. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1728–1738.
- Baraldi, L.; Grana, C.; and Cucchiara, R. 2015. A deep siamese network for scene detection in broadcast videos. In *Proceedings of the 23rd ACM international conference on Multimedia*, 1199–1202.
- Brown, A.; Kalogeiton, V.; and Zisserman, A. 2021. Face, body, voice: Video person-clustering with multiple modalities. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3184–3194.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9650–9660.
- Castro, S.; Hazarika, D.; Pérez-Rosas, V.; Zimmermann, R.; Mihalcea, R.; and Poria, S. 2019. Towards Multimodal Sarcasm Detection (An \_Obviously\_ Perfect Paper). In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4619–4629. Florence, Italy: Association for Computational Linguistics.
- Chen, L.; and Lee, C. M. 2017. Predicting audience’s laughter using convolutional neural network. *arXiv preprint arXiv:1702.02584*.
- Chen, S.; Nie, X.; Fan, D.; Zhang, D.; Bhat, V.; and Hamid, R. 2021. Shot contrastive self-supervised learning for scene boundary detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9796–9805.
- Dave, I.; Gupta, R.; Rizve, M. N.; and Shah, M. 2022. Tclr: Temporal contrastive learning for video representation. *Computer Vision and Image Understanding*, 219: 103406.
- Dorkenwald, M.; Xiao, F.; Brattoli, B.; Tighe, J.; and Modolo, D. 2022. Scvrl: Shuffled contrastive video representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4132–4141.
- Eysenck, H. J. 1942. The appreciation of humour: An experimental and theoretical study 1. *British Journal of Psychology. General Section*, 32(4): 295–309.
- Garner, R. 2005. Humor, analogy, and metaphor: HAM it up in teaching. *Radical Pedagogy*.
- Godfrey, J. J.; and Holliman, E. 1997. Switchboard-1 Release 2. *Linguistic Data Consortium, Philadelphia*, 926: 927.
- Hasan, M. K.; Rahman, W.; Zadeh, A. B.; Zhong, J.; Tanveer, M. I.; Morency, L.-P.; and Hoque, M. E. 2019. UR-FUNNY: A Multimodal Language Dataset for Understanding Humor. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2046–2056.
- Huang, Q.; Xiong, Y.; Rao, A.; Wang, J.; and Lin, D. 2020. Movienet: A holistic dataset for movie understanding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, 709–727. Springer.
- Jiang, Y.-G.; Jiang, Y.; and Wang, J. 2014. VCDB: a large-scale database for partial copy detection in videos. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, 357–371. Springer.
- Kayatani, Y.; Yang, Z.; Otani, M.; Garcia, N.; Chu, C.; Nakashima, Y.; and Takemura, H. 2021. The laughing machine: Predicting humor in video. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2073–2082.
- Kiddon, C.; and Brun, Y. 2011. That’s what she said: double entendre identification. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, 89–94.
- Kong, Q.; Cao, Y.; Iqbal, T.; Wang, Y.; Wang, W.; and Plumbley, M. D. 2020. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28: 2880–2894.
- Lei, J.; Li, L.; Zhou, L.; Gan, Z.; Berg, T. L.; Bansal, M.; and Liu, J. 2021. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7331–7341.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Lin, B.; Ye, Y.; Zhu, B.; Cui, J.; Ning, M.; Jin, P.; and Yuan, L. 2023. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*.
- Liu, Z.; Courant, R.; and Kalogeiton, V. 2022. Funnynet: Audiovisual learning of funny moments in videos. In *Proceedings of the Asian Conference on Computer Vision*, 3308–3325.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508: 293–304.
- Mun, J.; Shin, M.; Han, G.; Lee, S.; Ha, S.; Lee, J.; and Kim, E.-S. 2022. Bassl: Boundary-aware self-supervised learning for video scene segmentation. In *Proceedings of the Asian Conference on Computer Vision*, 4027–4043.

- Ni, B.; Peng, H.; Chen, M.; Zhang, S.; Meng, G.; Fu, J.; Xiang, S.; and Ling, H. 2022. Expanding language-image pretrained models for general video recognition. In *European Conference on Computer Vision*, 1–18. Springer.
- Patro, B. N.; Lunayach, M.; Srivastava, D.; Singh, H.; Nambodiri, V. P.; et al. 2021. Multimodal humor dataset: Predicting laughter tracks for sitcoms. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 576–585.
- Purandare, A.; and Litman, D. 2006. Humor: Prosody analysis and automatic recognition for f\* r\* i\* e\* n\* d\* s. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, 208–215.
- Qian, R.; Meng, T.; Gong, B.; Yang, M.-H.; Wang, H.; Belongie, S.; and Cui, Y. 2021. Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6964–6974.
- Rao, A.; Xu, L.; Xiong, Y.; Xu, G.; Huang, Q.; Zhou, B.; and Lin, D. 2020. A local-to-global approach to multimodal movie scene segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10146–10155.
- Rotman, D.; Porat, D.; and Ashour, G. 2016. Robust and efficient video scene detection using optimal sequential grouping. In *2016 IEEE international symposium on multimedia (ISM)*, 275–280. IEEE.
- Rui, Y.; Huang, T. S.; and Mehrotra, S. 1998. Exploring video structure beyond the shots. In *Proceedings. IEEE International Conference on Multimedia Computing and Systems (Cat. No. 98TB100241)*, 237–240. IEEE.
- Sarıgöz, Y. 2022. Triplet Loss - Advanced Intro. <https://qdrant.tech/articles/triplet-loss/>. Accessed: 2025-12-24.
- Schroff, F.; Kalenichenko, D.; and Philbin, J. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 815–823.
- Sikaroudi, M.; Ghogh, B.; Safarpour, A.; Karray, F.; Crowley, M.; and Tizhoosh, H. R. 2020. Offline versus online triplet mining based on extreme distances of histopathology patches. In *International Symposium on Visual Computing*, 333–345. Springer.
- Song, J.; Yang, Y.; Huang, Z.; Shen, H. T.; and Hong, R. 2011. Multiple feature hashing for real-time large scale near-duplicate video retrieval. In *Proceedings of the 19th ACM international conference on Multimedia*, 423–432.
- Souček, T.; and Lokoč, J. 2020. Transnet v2: An effective deep network architecture for fast shot transition detection. *arXiv preprint arXiv:2008.04838*.
- Stauffer, D. 1999. Let the good times roll: Building a fun culture. Harvard Management Update.
- Suls, J. M. 1972. A two-stage model for the appreciation of jokes and cartoons: An information-processing analysis. *The psychology of humor: Theoretical perspectives and empirical issues*, 1: 81–100.
- Thirard, P.-L.; and Codelli, L. 1994. Robert sklar. film, an international history of the medium, 1993; kristin thompson, david bordwell. film history, an introduction, 1994. *1895, revue d'histoire du cinema*, 17(1): 170–170.
- Tong, Z.; Song, Y.; Wang, J.; and Wang, L. 2022. Video-mae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35: 10078–10093.
- Vartabedian, R. A.; and Vartabedian, L. K. 1993. Humor in the Workplace: A Communication Challenge.
- Wang, L.; Xiong, Y.; Wang, Z.; Qiao, Y.; Lin, D.; Tang, X.; and Van Gool, L. 2016. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, 20–36. Springer.
- Wang, X.; Chen, H.; Wu, Z.; Zhu, W.; et al. 2024. Disentangled representation learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wanzer, M. B.; Frymier, A. B.; and Irwin, J. 2010. An explanation of the relationship between instructor humor and student learning: Instructional humor processing theory. *Communication education*, 59(1): 1–18.
- Weller, O.; and Seppi, K. 2019. Humor Detection: A Transformer Gets the Last Laugh. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3621–3625.
- Wu, X.; Hauptmann, A. G.; and Ngo, C.-W. 2007. Practical elimination of near-duplicates from web video search. In *Proceedings of the 15th ACM international conference on Multimedia*, 218–227.
- Xuan, H.; Stylianou, A.; and Pless, R. 2020. Improved embeddings with easy positive triplet mining. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2474–2482.
- Yang, D.; Lavie, A.; Dyer, C.; and Hovy, E. 2015. Humor recognition and humor anchor extraction. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, 2367–2376.
- Yu, Y.; Kim, J.; and Kim, G. 2018. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European conference on computer vision (ECCV)*, 471–487.
- Zhu, W.; Huang, Y.; Xie, X.; Liu, W.; Deng, J.; Zhang, D.; Wang, Z.; and Liu, J. 2023. AutoShot: A Short Video Dataset and State-of-the-Art Shot Boundary Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*.