

anomaly types, and temporal narratives, effectively eliminating the biases and limitations of Internet-collected datasets. Our pipeline integrates scene-conditioned anomaly assignment, multi-step storyline generation, and a temporally consistent long-form synthesis strategy that produces coherent 41-second videos with minimal human intervention. Extensive experiments demonstrate the scale, diversity, and complexity of Pistachio, revealing new challenges for existing methods and motivating future research on dynamic and multi-event anomaly understanding.

Keywords: Video Anomaly Detection · Video Anomaly Understanding · Synthetic Benchmark

1 Introduction

With the increasing deployment of autonomous driving, drone inspection, and large-scale surveillance systems, the ability to automatically detect abnormal events in videos has become indispensable. VAD [7, 23, 36, 53] aims to identify deviations from expected patterns, ranging from fires and traffic collisions to hazardous human activities. However, despite the steady growth of publicly available VAD datasets, existing benchmarks [1, 29, 41, 51] still fail to effectively and comprehensively evaluate modern VAD systems. Their limited diversity, biased anomaly distribution, and insufficient temporal complexity restrict both the reliability of evaluation and the generalization ability of trained models.

A closer examination reveals that these issues largely stem from how current datasets are sourced and curated. (i) Most videos are collected from the Internet, resulting in a heavy bias toward urban roadways and public streets. (ii) Anomaly types reflect events that most easily gain online visibility, producing severe long-tail distributions—frequent violent incidents but scarce environmental hazards, infrastructure failures, or rare accidents. (iii) The “normal” portions of the datasets are overly simplistic, typically featuring walking or light movement, which fails to capture the diversity of ordinary human or environmental activities. These weaknesses collectively cause VAD models to overfit to dataset-specific patterns, confuse complex-but-normal behaviors with anomalies, and fail to recognize rare or subtle events.

Beyond detection, the community has recently shifted toward VAU [12, 33, 48], which emphasizes deeper semantic and causal reasoning of identifying what occurred, why it occurred, and how the event unfolds over time. Yet current VAD benchmarks are intrinsically unable to evaluate these capabilities, because they lack structured temporal narratives, multi-step event progressions, and causal dependencies. Moreover, constructing a VAU benchmark is prohibitively expensive: it requires dense event-level annotations, temporal segmentation, and detailed semantic descriptions, especially when multiple anomalies co-occur within a single video. As a result, existing benchmarks remain limited in scope, failing to provide a comprehensive or scalable evaluation framework that fully captures the complexity of VAU.

In this paper, we propose *Pistachio*, a new VAD/VAU benchmark constructed entirely through a controlled, generation-based pipeline. Motivated by the remarkable progress of modern video generation models such as Sora [20], Veo 3 [42], and Wan [39], we argue that synthetic generation offers a more deterministic and scalable solution than filtering long-tailed Internet videos and performing labor-intensive manual annotation. Generation-based construction allows precise control over scenes, anomaly types, temporal progression, and visual diversity, providing balanced anomaly coverage and eliminating dataset biases that plague existing VAD resources.

Specifically, we propose several crucial designs to address the challenges in building a reliable, long-form anomaly benchmark. First, we develop a scene-conditioned anomaly assignment strategy that leverages Vision-Language Models (VLMs) to categorize source images into six major scene types and allocate contextually appropriate anomaly types, ensuring realistic and logically consistent pairings. Second, we design a multi-step storyline generation framework that decomposes each video into 7–8 descriptive segments for long videos (2–3 for short videos), enabling coherent narrative progression from normal activities to anomalous events. Third, we introduce a temporally consistent long-form video synthesis mechanism that chains short video clips using the last frame of each segment as the starting frame for the next, maintaining visual continuity across 41-second long videos (16-second for short videos). Finally, we implement a hybrid human-AI filtering approach that combines automated quality assessment with expert verification to ensure generated videos are free from artifacts and conform to real-world logic.

The resulting Pistachio-VAD benchmark is, to our knowledge, the most category-rich anomaly dataset to date, containing 4,962 videos spanning 1.68M frames, across 6 major scene categories and 31 diverse anomaly types, many of which do not appear in any prior work (e.g., landslides, animal predation, equipment breakdown). The dataset includes both static and moving cameras and features complex normal behaviors absent in existing benchmarks. In parallel, Pistachio-VAU contains 1,385 videos with event-level and video-level descriptions, including 35 videos with multiple co-occurring anomalies, which provides the first large-scale, fully annotated VAU benchmark constructed with zero manual annotation.

Our contributions can be summarized in three folds:

- We introduce Pistachio-VAD, a scalable, generation-based VAD benchmark that breaks scene and anomaly biases. It provides diverse scenes, balanced anomaly categories, and complex normal behaviors in large-scale, long-form videos.
- We design a fully automated pipeline for creating high-quality long-form anomaly videos. The controlled generation framework produces coherent 41-second videos through Scene-Aware Classification, Anomaly Type Specification, Multi-step storyline generation, and Temporally consistent long-form video synthesis, requiring minimal human intervention.

Table 1: Overview of existing VAD datasets.

Dataset	Year	Video	Frames	Scenes	Annotation	Detailed Anomaly types	Scalability
Subway Exit [3]	2008	1	38 940	1	Frame	3	×
Subway Entrance [3]	2008	1	86 535	1	Frame	5	×
UCSD Ped1 [16]	2010	70	14 000	1	Frame	5	×
UCSD Ped2 [16]	2010	28	4560	1	Frame	5	×
CUHK Avenue [21]	2013	35	30 652	1	Frame	5	×
ShanghaiTech [22]	2017	437	317 398	11	Frame	11	×
UCF-Crimes [32]	2018	1900	13 741 393	N/A	Frame	12	×
XD-Violence [43]	2020	4754	114 096	N/A	Frame	5	×
Street Scene [29]	2020	81	203 257	17	Frame	5	×
UBnormal [2]	2022	543	236 902	22	Pixel+Frame	22	×
NWPU Campus [5]	2023	547	1 466 073	43	Frame	28	×
MSAD [55]	2024	720	447 236	14	Frame	11	×
Pistachio (Ours)	2026	4962	1 676 822	3896	Frame+text	31	✓

- We construct Pistachio-VAU, the first large-scale VAU benchmark without manual annotation. We use generation storylines as structured semantic annotations, enabling comprehensive anomaly understanding including multi-event reasoning, without any manual labeling effort.

2 Related work

2.1 Video Anomaly Detection

VAD methods. Video Anomaly Detection methods [8, 10, 19, 27, 30, 31, 37, 44, 46, 47, 50, 54] can be categorized into semi-supervised [30, 50], weakly-supervised [8, 27, 31, 37, 54], and fully supervised [10, 19] methods. Recently, VLM- and LLM-based VAD has gained momentum. OVVD [44] synthesizes pseudo-unseen anomaly samples using a Novel Anomaly Synthesis module built on large generative vision models. Rule-based reasoning [46] converts normal frames into textual descriptions and uses LLMs to derive rules for identifying anomalies. Training-free pipelines [47] combine VLMs and LLMs for inference without model training. As VAD evolves toward richer contextual and causal reasoning, these text-driven approaches highlight the importance of semantic annotations: an aspect our benchmark directly supports.

VAD datasets. We report several statistics about the most utilized datasets in Table 1. As these statistics reveal, despite extensive research, existing datasets remain insufficient for evaluating real-world anomaly detection. Early benchmarks like CUHK Avenue [21], UCSD Ped2 [16], and Street Scene [29] focus on single-scene environments and simple human activities, limiting generalization. Even with more classes, Street Scene [29] suffers from severe long-tail imbalance (e.g., 61 jaywalking vs. 1 motorcycle-on-sidewalk). ShanghaiTech [22] expands the scene types but still includes only 158 anomaly instances across 11 categories. Larger datasets like UCF-Crime [32] and MASD [55] incorporate more realistic anomalies, yet their Internet- and surveillance-sourced videos inevitably reflect scene bias and overrepresented event types, failing to capture subtle or rare

anomalies. NWPU [5] Campus attempts to alleviate these issues through staged events performed by volunteers, but this approach is labor-intensive and still cannot cover the full diversity of abnormal behaviors. Overall, the scarcity and long-tail nature of real abnormal events make dataset construction inherently challenging. Our generation-based approach provides a scalable and controllable alternative, enabling balanced anomaly coverage, diverse scenes, and complex normal behaviors that are difficult to obtain from Internet-collected data.

2.2 Video Generation Models

Recent advances in text-to-video generation [15, 20, 39, 42] have dramatically improved the realism, controllability, and temporal coherence of synthetic videos. Diffusion- [24] and diffusion-transformer-based [25] models now support high-resolution, instruction-following video generation from text or images, often with strong physical plausibility. For example, OpenAI’s Sora [20] can generate minute-scale, photorealistic videos directly from prompts and input images, maintaining complex camera motion and scene dynamics. Google’s Veo series [42] similarly produces high-quality clips with improved modeling of real-world physics and human motion, targeting cinematic applications. In parallel, open and industrial models such as Wan [39] adopt diffusion transformer architectures and large-scale pre-training to deliver efficient 720p, 24fps video generation for both text-to-video and image-to-video settings. These systems demonstrate strong visual fidelity over short durations (typically 5–10 seconds), but naively extending them to long videos often leads to temporal drift, artifacts, or prompt forgetting. Our work builds on this line of research by introducing a long-form generation mechanism that composes multiple short clips into coherent 41-second narratives.

3 Pistachio Benchmark

3.1 Dataset Statistics

Pistachio for Video Anomaly Detection We introduce our new dataset for VAD, a comprehensive collection comprising 4,962 videos, which total 29.11 hours of footage captured at 16 FPS. It is the largest known semi-supervised VAD dataset. The dataset is structured across 6 major scenes and encompasses 31 distinct anomaly categories. Notably, we introduce 10 novel anomaly types that are absent in existing VAD benchmarks: Animal Predation, Animal Abuse, Construction Accidents, Infrastructure Failure, Natural Disasters, Avalanche, Equipment Breakdown, Extreme Weather Events, Ground Collapse, and Landslide. In addition to frame-level labels, it provides rich textual annotations. The detailed composition of our training and testing sets is provided in the supplementary material.

To highlight the contributions of our dataset, we comprehensively compare it to existing benchmarks. Our dataset is distinguished by several outstanding

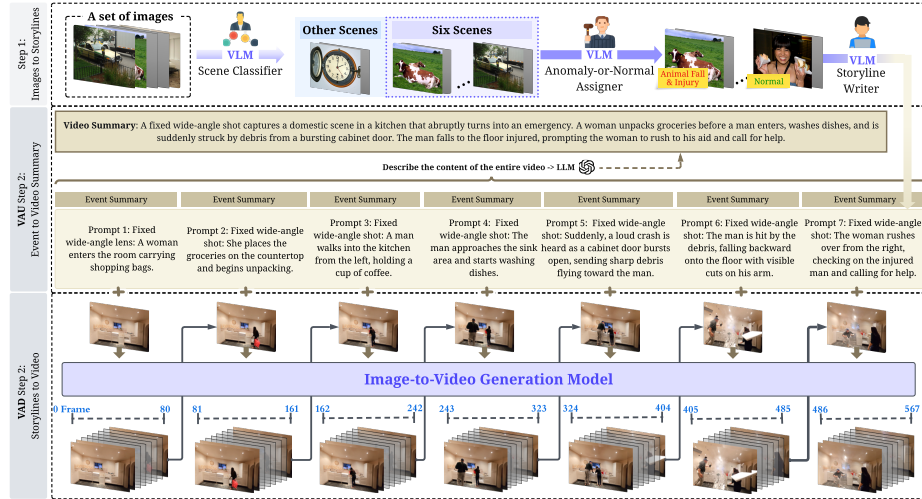


Fig. 2: Overview of our video anomaly dataset generation pipeline. **Step 1 (Storyline Generation):** A VLM-based scene classifier processes input images to identify scenes. The Anomaly Type Allocator then assigns appropriate anomaly types (e.g., Fire) to generate coherent storylines across six scenes by LLMs. **VAU Step 2 (Event-to-Video Summary):** The storyline is decomposed into event summaries, with each event described by detailed prompts specifying camera angles, actions, and temporal progression. **VAD Step 2 (Storyline-to-Video):** An image-to-video generation model synthesizes coherent video sequences from the event summaries, producing temporally consistent frames that maintain narrative continuity across the entire storyline.

traits. First, as illustrated in 3, it introduces a wealth of scene-specific anomalies, like Avalanche, Landslide, and Animal Predation. And our dataset integrates novel categories—nearly half of which are absent in other datasets—to emphasize critical social and ecological safety issues. We have maintained a balanced categorical distribution, directly addressing the persistent long-tail problem. Second, the dataset’s unprecedented scale in scene diversity comprehensively mitigates the long-standing issue of model scene-dependency. Third, as shown in the examples in Fig 1 our normal samples extend beyond trivial actions; thanks to a meticulous storyline design, we incorporate a rich array of complex daily activities, like handshaking and vehicle boarding. Fourth, Finally, the inclusion of both conventional static shots to mimic surveillance perspectives and dynamic mobile camera shots makes our dataset uniquely suited for emerging applications in fields such as drone monitoring and autonomous driving.

Pistachio for Video Anomaly Understanding For the task of Video Anomaly Understanding, we have curated a specialized dataset comprising 1,385 videos, which total 517,514 frames. A key innovation of this dataset is its introduction of multiple, distinct anomaly types within a single video—a first for VAU bench-

marks. Furthermore, each video is supported by comprehensive event-level and video-level annotations. Echoing the strengths of its VAD counterpart, our VAU dataset also stands as the most extensive benchmark to date in terms of scene richness, anomaly diversity, and balanced categorical distribution.

3.2 Multi-Stage Data Generation Pipeline

Our multi-stage data generation pipeline is meticulously designed to create a large-scale, high-quality, and richly annotated dataset. The entire process is divided into three key stages.

First, Video Data Generation Sec 3.2 details the core synthesis pipeline. As shown in Fig. 2, this process involves three key steps: Step 1 uses a VLM to generate scene-conditioned anomaly storylines from static images, Step 2 decomposes these storylines into detailed event-level prompts, and Step 3 employs a video generation model to synthesize these into the final video clips. This is followed by Video Data Filtering Sec 3.2, a crucial quality control step that employs a hybrid human-AI screening approach to ensure the realism and logical consistency of all generated videos. Finally, and specifically to support the VAU task, the Annotation Generation stage Sec 3.2 leverages an LLM/VLM to automatically process the storylines, generating the rich, multi-granularity (event-level and video-level) annotations.

Video Data Generation

Scene-Aware Classification. We utilize images $\mathcal{I} = \{I_1, \dots, I_N\}$ sourced from the COCO 2017 dataset. A Vision-Language Model $\mathcal{M}(\cdot)$ (VLM) categorizes each image I_i into one of $K = 6$ predefined scene categories extracted from scene configuration $\mathcal{C}_{scene}$, plus an additional “Other” category, yielding $C = \{c_1, \dots, c_K\} \cup \{\text{“Other”}\}$. The categorization prompt is provided in the supplementary materials.

This stage assigns each image to a scene group:

$$I_i \rightarrow \mathcal{G}_{\hat{c}_i}, \text{ where } \hat{c}_i = \mathcal{M}(I_i, \text{prompt}_{scene}) \quad (1)$$

where $\mathcal{G}_{c_j}$ denotes the set of images assigned to scene category c_j .

Anomaly Type Specification. For each scene category $c_j \in C$, we manually define a tailored set of plausible anomaly types $A_j = \{a_{j,1}, \dots, a_{j,M_j}\}$ in $\mathcal{C}_{scene}[c_j]$ (detailed in supplementary materials). The VLM assigns specific anomaly types to each image within the scene group through an anomaly mapping $\mathcal{A} : I_i \mapsto a_i$:

$$\hat{a}_i = \mathcal{M}(I_i, \phi_j), \text{ where } \phi_j = \text{BuildPrompt}(c_j, A_j) \quad (2)$$

The assignment prompt templates are provided in the supplementary materials.

Multi-step storyline generation. For each image I_i with assigned scene c_i and anomaly type a_i , we retrieve the corresponding prompt template $p_i = \mathcal{C}_{scene}[c_i][a_i]$ and format it to generate a coherent video storyline:

$$\mathcal{S}_i = \mathcal{M}(I_i, \psi_i), \text{ where } \psi_i = \text{FormatPrompt}(p_i, I_i) \quad (3)$$

Each storyline $\mathcal{S}_i = \{s_1, \dots, s_L\}$ comprises L descriptive segments, where $L \in [7, 8]$ for long videos and $L \in [2, 3]$ for short videos. Detailed prompt templates are in the supplementary materials.

Temporally consistent long-form video synthesis. The final stage synthesizes video clips from the generated storylines. For each storyline $\mathcal{S}_i$, its segments $\{s_1, s_2, \dots, s_L\}$ are sequentially fed into the Wan video generation model. For each segment s_l , a video clip $V_l = \text{Wan}(I_{l-1}, s_l)$ is generated, where I_{l-1} is the starting frame. The initial frame I_0 is the original image I_i from $\mathcal{I}$, and for subsequent segments ($l > 1$), the last frame of V_{l-1} serves as I_{l-1} to ensure temporal continuity. All clips $\{V_1, V_2, \dots, V_L\}$ are concatenated to form the final video V_{final}^i .

Video Data Filtering Balanced and highly realistic AI-generated data is essential for video anomaly detection. Our video filtering process consists of two stages: first, we employ VideoScore [13] to automatically score all generated videos across multiple quality dimensions, discarding those in the bottom 5% (average score below 3.98) as a preliminary quality control step. Subsequently, we conduct rigorous manual filtering to ensure the remaining videos meet the following criteria:

1. The anomaly events in the videos are reasonably designed and conform to real-life logic.
2. No obvious signs of AI generation, such as object inconsistencies or sudden visual distortions.
3. The camera angles, viewpoints, and lighting conditions are realistic and emulate those of genuine videos.

This rigorous two-stage filtering ensures high visual fidelity and physical consistency across the dataset, guaranteeing that the final videos are free from generative artifacts and suitable for robust anomaly analysis.

Annotation Generation and Refinement To ensure maximum rigor and accuracy, we adopt a hybrid annotation strategy. For VAD task, we utilize manual human annotation to define ground truth labels. For VAU task, we implement a multi-stage pipeline to generate event-level and video-level text annotations.

Our annotation workflow originates from the initial video storylines, which provide inherently precise, localized event-level descriptions. We utilize these texts both as the input prompts for video generation and as the foundational ground truth for anomaly understanding. To generate holistic video-level descriptions, we utilize LLMs to summarize individual event annotations into a

single, coherent video summary, as illustrated in Fig. 2, Step 2. This process accounts for temporal partitions and contextual variances within the storyline.

To further enhance the fidelity of the generated text, we introduce a VLM Refiner. The VLM performs a secondary pass, cross-referencing the generated summaries with the actual visual content to correct inaccuracies or misalignments. This automated yet refined workflow enables the production of high-quality, multi-granularity labels, ensuring both consistency and scalability across the dataset without the exhaustive cost of manual text writing.




Fig. 3: Distribution of anomaly videos across different anomaly types. For each type, the left bar represents short videos and the right bar represents long videos. The right-most single bars indicate multi-anomaly videos, which are exclusively long-form.

4 Experiments

4.1 Dataset Construction

To facilitate a comprehensive evaluation of long-form video generation, we constructed a new, diverse, multi-source dataset. The foundational imagery is derived from the **COCO** dataset [18], which provides a rich repository of common objects and everyday scenes.

However, we identified significant gaps in COCO’s scene coverage. To remedy this, we augmented the dataset with two supplementary sources: (1) **MIT Indoor Scenes** [28], to enrich the diversity of indoor environments; and (2) **Web-sourced Images**, a curated collection of high-resolution images depicting large-scale infrastructures, such as bridges and tunnels, which are notably scarce in existing benchmarks.

The complete dataset generation process (including pipeline-based video synthesis, described below) was computationally intensive, requiring approximately

Table 2: Performance comparison on the Pistachio dataset for VAD. All methods are trained on the Pistachio training set. All metrics are reported in percentage (%).

Method	Year	Backbone	Public Safety		Accidents Infrastructure		Natural Hazards		Health Medical		Animal Incidents		Overall	
			AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP
RTFM [38]	2021	I3D	69.0	62.3	76.3	84.3	78.2	83.1	68.9	69.1	76.6	75.9	82.9	69.3
DR-DMU [53]	2023	I3D	66.5	63.4	79.8	88.4	75.1	83.1	76.0	78.8	69.8	64.0	81.5	71.5
MGFN [7]	2023	I3D	63.1	57.2	69.9	74.6	74.5	82.4	66.6	65.4	54.1	52.2	74.9	50.2
CLIP-TSA [14]	2023	I3D	66.9	58.5	78.7	83.3	76.4	82.9	69.9	62.1	64.5	56.5	76.9	55.2
CLIP-TSA [14]	2023	ViT	60.8	56.0	57.7	65.2	47.6	55.4	54.1	48.6	69.2	59.7	80.9	57.3
MULDE [23]	2024	I3D	50.1	40.1	62.0	58.4	52.0	49.2	57.5	37.8	55.1	43.9	63.4	34.9
VadCLIP [45]	2024	ViT	63.0	56.7	71.0	80.9	56.5	66.5	67.3	66.2	66.4	62.3	78.1	64.1
PELVAD [26]	2024	ViT	71.2	65.3	72.1	81.6	59.4	65.1	66.5	65.2	70.9	71.5	83.7	71.0
Fed-WSVAD [40]	2025	ViT	62.2	54.8	79.7	86.1	74.6	83.0	71.2	64.6	49.3	47.3	83.2	71.9
VADTree [17]	2025	—	64.99	44.21	63.37	42.15	53.73	62.76	68.09	52.74	57.30	39.00	72.52	25.15

20 days of processing time on a cluster of 32 NVIDIA A100 (80GB) GPUs. For the manual filtering stage, the entire process required approximately 2 person-days, with acceptance rates of 90% for normal videos and 80% for anomaly videos. This composite benchmark enables a more challenging and thorough assessment of a model’s generalization capabilities and scene understanding.

4.2 Implementation Details

Our proposed hierarchical generation pipeline (detailed in Section 3.2) employs several foundation models as its backbone:

- **Video Generation Model:** For the core synthesis task, we employ the Wan [39] video diffusion model as our T2V generator, valued for its high-fidelity output and temporal coherence.
- **Image-to-Storyline VLM:** We utilize Qwen2.5-VL-32B-Instruct [4] as our primary VLM. It is responsible for parsing the initial frame and generating a structured, high-level "storyline" or sequence of events.
- **Event-to-Video-Level Model:** To bridge the fine-grained to holistic semantic gap, Qwen3-8B [34] acts as a semantic aggregator, synthesizing discrete event annotations into coherent video summaries. Subsequently, Qwen3-VL-32B [4] serves as a VLM Refiner, cross-referencing these summaries with visual content to rectify any temporal or factual misalignments.

4.3 Analysis of Alternative Pipelines

To motivate our design, we first investigated several intuitive baseline pipelines for long-form video generation and identified their critical failure modes.

1. **Direct Long-Form Generation:** This one-shot approach tasks the model with generating the entire video from a single prompt. This method consistently suffers from two issues: (a) **Ghosting and Artifacts:** The model fails to maintain object identity and scene consistency over long durations. (b) **Prompt Forgetting:** The generated content progressively diverges from the initial prompt, indicating a severe semantic drift.

2. **Fixed First-Last Frame (Looping):** This pipeline uses a first-last-frame-to-video model, setting the initial and final frames to be identical for seamless looping. This constraint, however, induces pathological motion artifacts. The model appears to “rush” to satisfy the end-frame constraint, resulting in unnatural accelerations, “stuttering” movements, and physically implausible object trajectories as it attempts to force the scene back to its origin.
3. **Naive Storyline Chaining:** This method generates video in short segments. For each new segment, the prompt is slightly modified based on the content of the previous segment’s first frame. We found this approach lacks precise control. The model interprets the modified prompt as a simple “continuation” of the prior scene’s state, rather than as an instruction to execute a new, distinct “event” as intended.

4.4 Video Anomaly Detection

We design three comprehensive sets of experiments to investigate the characteristics and utility of the Pistachio dataset. First, To evaluate the challenge our newly proposed Pistachio dataset poses to existing VAD methods, we selected several mainstream baseline methods for assessment. Second, we conduct cross-dataset generalization experiments to examine the adaptability of models to new scenarios and perspectives. Moreover, to ensure our findings are not restricted by potential biases inherent to the Wan generation model, we incorporate additional evaluation sets generated by Hailuo and Cosmos. Finally, acknowledging that Pistachio features diverse viewpoints and cinematic styles—distinct from traditional fixed-angle surveillance datasets—we investigate the impact of data scaling and domain fusion by combining Pistachio with existing benchmarks to observe performance gains on real-world tasks.

Implementation Details. We selected nine widely used VAD methods as baselines: RTFM [38], DR-DMU [53], MGFN [7], MULDE [23], CLIP-TSA [14], VadCLIP [45], PEL4VAD [26], Fed-WSVAD [40], and VADTree [17]. Following standard evaluation protocols, the training-based methods are trained on the Pistachio training set and evaluated on our test set, while the training-free method VADTree is directly deployed for evaluation without any parameter updates. Performance is measured using frame-level AUC and AP. To prepare the data for the training-based methods, we extract frame-level features using pre-trained visual models, specifically I3D [6] and ViT [11], following the processing approaches in [38] and [45]. For VADTree, we directly process the raw videos by leveraging pre-trained Generic Event Boundary Detection (GEBD) and video-language models following its original implementation.

Performance Comparison. As shown in Table 2, we report the performance of baseline methods on the Pistachio test set. Generally, most baseline methods exhibit suboptimal performance in the **Public Safety** and **Animal Incidents** categories. This performance gap suggests that current VAD models still struggle to capture fine-grained behavioral features and complex biological motion.

Table 3: Cross-dataset generalization performance across five target benchmarks and MSAD. Models are evaluated on both synthetic (Pistachio, Cosmos, Hailuo) and real-world (UCF-Crime, XD-Violence, MSAD) datasets. **Joint** denotes training on the union of all five datasets and is listed for reference. Best results among single-source training are in **bold**. AUC/AP metrics are reported in %.

Train Set	Method	Pistachio		Cosmos		Hailuo		UCF-Crime		XD-Violence		MSAD	
		AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP
Pistachio	Fed-WSVAD	83.23	71.86	81.09	78.76	83.97	79.46	55.69	10.14	70.11	44.51	81.68	24.28
	CLIP-TSA	80.91	57.30	72.17	58.74	75.89	59.12	61.58	10.48	86.21	62.48	77.12	54.24
UCF-Crime	Fed-WSVAD	71.46	58.69	72.29	78.37	80.82	72.24	84.35	24.49	—	—	—	—
	CLIP-TSA	61.06	39.80	71.54	58.66	68.71	70.84	82.56	21.62	81.93	51.75	—	—
XD-Violence	Fed-WSVAD	71.61	62.14	81.67	81.43	83.36	71.64	78.65	19.64	92.50	77.04	86.44	32.76
	CLIP-TSA	68.05	42.49	73.46	68.74	74.18	56.90	80.11	23.36	92.58	74.18	79.35	18.69
Joint	CLIP-TSA	82.17	67.87	83.04	87.37	87.58	83.60	85.89	28.36	93.53	80.46	85.73	25.69

Notably, while previous benchmarks often overlooked animal-related anomalies, the Pistachio dataset introduces these challenging scenarios, revealing the limitations of existing motion-based representations.

Among the evaluated methods, **PEL4VAD** [26] and our **Fed-WSVAD** [40] achieve superior results, which can be attributed to their distinct yet complementary strengths. Commonly, both methods leverage vision-language association (e.g., CLIP-based features), allowing the models to integrate high-level semantic priors to disambiguate complex visual scenes. Individually, PEL4VAD excels by utilizing a prompt-enhanced learning mechanism to aggregate temporal context, which is particularly effective for categories requiring long-range dependency modeling. In contrast, our Fed-WSVAD benefits from a federated learning framework that captures a broader distribution of anomalies from decentralized clients. By aggregating diverse local knowledge into a global context-driven model, Fed-WSVAD achieves a more balanced and robust representation, leading to the highest Overall AP of 71.9%. The challenges posed by our dataset’s enriched normal diversity are particularly evident when examining methods that lack semantic guidance. The poor performance of MULDE (63.4% AUC, 34.9% AP) stems from its likelihood-based design, which struggles to define a tight boundary over highly variable normal data. More notably, the training-free VADTree yields a catastrophic AP drop to 25.15% (72.52% AUC); without in-domain training, its zero-shot vision-language reasoner is highly vulnerable to our enriched normal diversity, leading to severe false positives. These results underscore that neither traditional likelihood models nor zero-shot training-free approaches can robustly handle the complex normal variations and interactive anomalies introduced by Pistachio.

Cross-Dataset Generalization and Bias Analysis As shown in Table 3, the Pistachio dataset demonstrates strong cross-dataset generalization across diverse benchmarks, particularly on the synthetic Cosmos and Hailuo benchmarks as well as the real-world XD-Violence and MSAD datasets. Notably, the com-

Table 4: Impact of training data scaling on UCF-Crime test performance for CLIP-TSA. We evaluate performance by combining our Pistachio dataset with varying proportions of the UCF-Crime training set. All results are reported as AUC / AP (%).

Pistachio		UCF		Pist.+10% UCF		Pist.+30% UCF		Pist.+50% UCF	
AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP
61.58	10.48	82.56	21.62	83.96	24.25	86.19	28.44	86.79	33.04

petitive—and in some cases superior—performance of Pistachio-trained models compared to both UCF-Crime- and XD-Violence-trained models demonstrates promising transferability of our synthetic samples. This outperformance can be attributed to Pistachio’s greater scene diversity and balanced anomaly distribution, which together provide richer and more generalizable representations, underscoring its practical reliability as a training source for diverse real-world anomaly scenarios.

However, we observe that models trained solely on Pistachio show lower performance on UCF-Crime, due to the latter’s low resolution and fixed-angle surveillance perspectives, as well as the inherent gap between synthetic and real-world data. To enhance portability to legacy surveillance scenarios, we explore a data scaling strategy. As illustrated in Table 4, even with minimal fine-tuning using a small fraction of UCF-Crime data, the Pistachio-pretrained model achieves significant gains. Specifically, integrating Pistachio with just 10% of the UCF training set outperforms a model trained on the full UCF dataset, demonstrating that Pistachio serves as a powerful foundational representation that complements traditional domain-specific data. Finally, the joint training results in Table 3 further confirm the complementary nature of synthetic and real-world data: combining all datasets consistently achieves the strongest performance across all benchmarks, suggesting that Pistachio and real-world data are most powerful when used together. However, we also acknowledge that Pistachio, as a synthetic dataset, carries inherent limitations in fully substituting real-world data, which we discuss in detail in Section 5.

Finally, to rule out generator-specific artifacts, we evaluate two dedicated VAD benchmarks built with Cosmos-2.5 and Hailuo-2.3. The Pistachio-trained models achieve consistent and significant improvements, raising the AUC from 71.54% to 85.22% on Cosmos and from 68.71% to 82.34% on Hailuo. These cross-generator enhancements confirm that Pistachio offers generalizable anomaly representations rather than overfitting to specific generator biases, establishing the robustness and diversity of our dataset.

4.5 Video Anomaly Understanding

To evaluate the benchmarks of video anomaly understanding, we adopt F1-Score as the main metric to measure the model’s capability in comprehending anomalies across different temporal granularities. The F1-Score is calculated as

Table 5: Experiments of Pistachio datasets on Visual Anomaly Understanding, spanning multiple model families and scales. We differentiate understanding at the event-level and video-level outputs.

Model	Params	Event-level $\uparrow$	Video-level $\uparrow$	Avg. $\uparrow$
InternVL3-1B	1B	28.25	23.04	25.65
InternVL3-2B	2B	30.72	25.22	27.97
Qwen2.5-VL-3B	3B	29.24	23.91	26.58
LLaVA-Next-Video-7B	7B	29.61	23.16	26.39
Qwen2.5-VL-7B	7B	30.57	25.58	28.08
Qwen3-VL-8B	8B	28.14	26.65	27.40
InternVL3-8B	8B	32.56	25.65	29.11
InternVL3-14B	14B	32.29	26.49	29.39

the harmonic mean of precision and recall, providing a balanced assessment of the model’s performance.

Evaluation Protocol. Following the hierarchical structure of our HIVA-70k [49] dataset, we assess the model’s understanding ability at two key levels: 1) **Event-level**, where the model analyzes anomaly events spanning multiple clips, requiring intermediate-term reasoning capabilities, and 2) **Video-level**, where the model comprehends the entire video narrative, demanding long-term contextual understanding. For each level, we measure the F1-Score by comparing generated descriptions with ground truth annotations, counting a prediction as correct if semantic similarity exceeds a predefined threshold.

Comparison with State-of-the-Art Models. As shown in Tab. 5, we compare against several state-of-the-art Multimodal Large Language Models (MLLMs) spanning two families and a wide range of scales, including InternVL3 (1B–14B) [9], Qwen2.5-VL [4], Qwen3-VL-8B [35], and LLaVA-Next-Video-7B [52]. The results reveal significant variations across model sizes and architectures, confirming that performance on Pistachio-VAU is sensitive to both model family and parameter scale. InternVL3-14B attains the best average score, while InternVL3-8B leads at the event level and Qwen3-VL-8B at the video level.

Analysis. The performance gap between event-level and video-level understanding reveals that models exhibit stronger capabilities in analyzing shorter temporal segments than comprehending extended video narratives. This observation supports our motivation for developing hierarchical instruction data, highlighting the challenge of maintaining contextual coherence over longer sequences. Furthermore, larger model parameters do not always guarantee better performance. For instance, the 8B-parameter Qwen3-VL achieves results comparable to 7B models, and within the InternVL3 family the gains from 8B to 14B are marginal, suggesting that model architecture and training data quality play crucial roles beyond raw parameter count.

5 Scope and Limitations

The most fundamental limitation of Pistachio is the domain gap between synthetic and real-world data, which we characterise here openly. A model trained

only on Pistachio transfers to real-world datasets with variable generalization performance, exposing a clear cross-domain discrepancy. On legacy, low-resolution surveillance data like UCF-Crime, CLIP-TSA reaches 61.58% AUC / 10.48% AP, dropping below the 82.56% / 21.62% obtained from training on UCF-Crime itself (Tables 3 and 4). Although this performance variance is dataset-dependent—with Pistachio exhibiting relatively better transferability to MSAD and XD-Violence—the pronounced degradation observed on UCF-Crime demonstrates that models trained solely on our synthetic benchmark still encounter generalization limitations when facing specific real-world data distributions.

To help the community better navigate this gap, we frankly outline the conditions under which Pistachio can and cannot be expected to substitute for real-world data. Specifically, Pistachio **cannot** serve as a direct substitute when the target task is exclusively restricted to traditional fixed-camera surveillance, or when it heavily relies on long-term character identity consistency. Despite our rigorous manual filtering, current video generation foundations still occasionally introduce subtle appearance shifts in characters over time. Conversely, Pistachio **can** effectively substitute for real-world data when the primary objective is to evaluate and enhance a model’s capacity for detecting sudden, unexpected anomalies across generalized scenarios. For these tasks, Pistachio’s rich anomaly categories, highly diverse storylines, and strictly balanced distribution provide a robust, bias-free evaluation bed that real-world benchmarks cannot replicate.

We accordingly offer concrete guidance for future users: Pistachio is best utilized as a scalable, high-diversity training foundation and a comprehensive gauge for generalized anomaly coverage. We recommend practitioners treat Pistachio-only metrics as a robust benchmark for broad anomaly comprehension, and perform modest in-domain fine-tuning when deploying to surveillance settings with low resolution or high-angle viewpoints.

6 Conclusion

In this paper, we introduce Pistachio, a new large-scale, Synthetic, Balanced, and Long-Form benchmark for VAD and VAU, accompanied by a highly automated and publicly released generation pipeline. Our comprehensive evaluation demonstrates that Pistachio’s unique characteristics pose significant challenges to existing models, exposing their limitations. Furthermore, our synthetic approach provides precise ground truth for VAU, mitigating the heavy reliance on laborious manual annotation. We also provide concrete guidance for practitioners, recommending Pistachio as a scalable complement to real-world data and a foundation for fine-tuning rather than an unconditional substitute, and we hope it will serve as a valuable resource for future VAD and VAU research.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant 62506235.

References

1. Acintoae, A., Florescu, A., Georgescu, M., Mare, T., Sumedrea, P., Ionescu, R.T., Khan, F.S., Shah, M.: Ubnormal: New benchmark for supervised open-set video anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2022)
2. Acintoae, A., Florescu, A., Georgescu, M.I., Mare, T., Sumedrea, P., Ionescu, R.T., Khan, F.S., Shah, M.: Ubnormal: New benchmark for supervised open-set video anomaly detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20143–20153 (2022)
3. Adam, A., Rivlin, E., Shimshoni, I., Reinitz, D.: Robust real-time unusual event detection using multiple fixed-location monitors. *IEEE transactions on pattern analysis and machine intelligence* **30**(3), 555–560 (2008)
4. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., Zhong, H., Zhu, Y., Yang, M., Li, Z., Wan, J., Wang, P., Ding, W., Fu, Z., Xu, Y., Ye, J., Zhang, X., Xie, T., Cheng, Z., Zhang, H., Yang, Z., Xu, H., Lin, J.: Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923* (2025)
5. Cao, C., Lu, Y., Wang, P., Zhang, Y.: A new comprehensive benchmark for semi-supervised video anomaly detection and anticipation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20392–20401 (2023)
6. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
7. Chen, Y., Liu, Z., Zhang, B., Fok, W., Qi, X., Wu, Y.C.: Mgn: Magnitude-contrastive glance-and-focus network for weakly-supervised video anomaly detection (2022)
8. Chen, Y., Liu, Z., Zhang, B., Fok, W., Qi, X., Wu, Y.C.: Mgn: Magnitude-contrastive glance-and-focus network for weakly-supervised video anomaly detection. In: Proceedings of the AAAI conference on artificial intelligence. vol. 37, pp. 387–395 (2023)
9. Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al.: Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24185–24198 (2024)
10. Del Giorno, A., Bagnell, J.A., Hebert, M.: A discriminative framework for anomaly detection in large videos. In: European conference on computer vision. pp. 334–349. Springer (2016)
11. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale (2021), <https://arxiv.org/abs/2010.11929>
12. Du, H., Zhang, S., Xie, B., Nan, G., Zhang, J., Xu, J., Liu, H., Leng, S., Liu, J., Fan, H., Huang, D., Feng, J., Chen, L., Zhang, C., Li, X., Zhang, H., Chen, J., Cui, Q., Tao, X.: Uncovering what, why and how: A comprehensive benchmark for causation understanding of video anomaly (2024), <https://arxiv.org/abs/2405.00181>
13. He, X., Jiang, D., Zhang, G., Ku, M., Soni, A., Siu, S., Chen, H., Chandra, A., Jiang, Z., Arulraj, A., Wang, K., Do, Q.D., Ni, Y., Lyu, B., Narsupalli, Y., Fan, R., Lyu, Z., Lin, Y., Chen, W.: Videoscore: Building automatic metrics to simulate fine-grained human feedback for video generation. *ArXiv abs/2406.15252* (2024), <https://arxiv.org/abs/2406.15252>

14. Joo, H.K., Vo, K., Yamazaki, K., Le, N.: Clip-tsa: Clip-assisted temporal self-attention for weakly-supervised video anomaly detection (2023), <https://arxiv.org/abs/2212.05136>
15. Kong, W., Tian, Q., Zhang, Z., Min, R., Dai, Z., Zhou, J., Xiong, J., Li, X., Wu, B., Zhang, J., Wu, K., Lin, Q., Yuan, J., Long, Y., Wang, A., Wang, A., Li, C., Huang, D., Yang, F., Tan, H., Wang, H., Song, J., Bai, J., Wu, J., Xue, J., Wang, J., Wang, K., Liu, M., Li, P., Li, S., Wang, W., Yu, W., Deng, X., Li, Y., Chen, Y., Cui, Y., Peng, Y., Yu, Z., He, Z., Xu, Z., Zhou, Z., Xu, Z., Tao, Y., Lu, Q., Liu, S., Zhou, D., Wang, H., Yang, Y., Wang, D., Liu, Y., Jiang, J., Zhong, C.: Hunyuanvideo: A systematic framework for large video generative models (2025), <https://arxiv.org/abs/2412.03603>
16. Li, W., Mahadevan, V., Vasconcelos, N.: Anomaly detection and localization in crowded scenes. *IEEE transactions on pattern analysis and machine intelligence* **36**(1), 18–32 (2013)
17. Li, W., Xu, Y., Rao, Y., Wang, Z., Deng, S.: Vadtree: Explainable training-free video anomaly detection via hierarchical granularity-aware tree (2025), <https://arxiv.org/abs/2510.22693>
18. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer (2014)
19. Liu, K., Ma, H.: Exploring background-bias for anomaly detection in surveillance videos. In: *Proceedings of the 27th ACM International Conference on Multimedia*. pp. 1490–1499 (2019)
20. Liu, Y., Zhang, K., Li, Y., Yan, Z., Gao, C., Chen, R., Yuan, Z., Huang, Y., Sun, H., Gao, J., et al.: Sora: A review on background, technology, limitations, and opportunities of large vision models. *arXiv preprint arXiv:2402.17177* (2024)
21. Lu, C., Shi, J., Jia, J.: Abnormal event detection at 150 fps in matlab (2013)
22. Luo, W., Liu, W., Gao, S.: A revisit of sparse coding based anomaly detection in stacked rnn framework. In: *Proceedings of the IEEE international conference on computer vision*. pp. 341–349 (2017)
23. Micorek, J., Possegger, H., Narnhofer, D., Bischof, H., Kozinski, M.: Mulde: Multiscale log-density estimation via denoising score matching for video anomaly detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 18868–18877 (June 2024)
24. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4195–4205 (2023)
25. Peebles, W., Xie, S.: Scalable diffusion models with transformers (2023), <https://arxiv.org/abs/2212.09748>
26. Pu, Y., Wu, X., Wang, S.: Learning prompt-enhanced context features for weakly-supervised video anomaly detection. *arXiv preprint arXiv:2306.14451* (2023)
27. Pu, Y., Wu, X., Yang, L., Wang, S.: Learning prompt-enhanced context features for weakly-supervised video anomaly detection. *IEEE Transactions on Image Processing* (2024)
28. Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 413–420. IEEE (2009)
29. Ramachandra, B., Jones, M.: Street scene: A new dataset and evaluation protocol for video anomaly detection. In: *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. pp. 2569–2578 (2020)
30. Ristea, N.C., Croitoru, F.A., Ionescu, R.T., Popescu, M., Khan, F.S., Shah, M., et al.: Self-distilled masked auto-encoders are efficient video anomaly detectors. In:

- Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 15984–15995 (2024)
31. Sultani, W., Chen, C., Shah, M.: Real-world anomaly detection in surveillance videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6479–6488 (2018)
 32. Sultani, W., Chen, C., Shah, M.: Real-world anomaly detection in surveillance videos (2019), <https://arxiv.org/abs/1801.04264>
 33. Tang, J., Lu, H., Wu, R., Xu, X., Ma, K., Fang, C., Guo, B., Lu, J., Chen, Q., Chen, Y.C.: Hawk: Learning to understand open-world video anomalies (2024), <https://arxiv.org/abs/2405.16886>
 34. Team, Q.: Qwen3 (April 2025), <https://qwenlm.github.io/blog/qwen3/>
 35. Team, Q.: Qwen3 technical report (2025), <https://arxiv.org/abs/2505.09388>
 36. Tian, Y., Pang, G., Chen, Y., Singh, R., Verjans, J.W., Carneiro, G.: Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. arXiv preprint arXiv:2101.10030 (2021)
 37. Tian, Y., Pang, G., Chen, Y., Singh, R., Verjans, J.W., Carneiro, G.: Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4975–4986 (2021)
 38. Tian, Y., Pang, G., Chen, Y., Singh, R., Verjans, J.W., Carneiro, G.: Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. arXiv preprint arXiv:2101.10030 (2021)
 39. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., Zeng, J., Wang, J., Zhang, J., Zhou, J., Wang, J., Chen, J., Zhu, K., Zhao, K., Yan, K., Huang, L., Feng, M., Zhang, N., Li, P., Wu, P., Chu, R., Feng, R., Zhang, S., Sun, S., Fang, T., Wang, T., Gui, T., Weng, T., Shen, T., Lin, W., Wang, W., Wang, W., Zhou, W., Wang, W., Shen, W., Yu, W., Shi, X., Huang, X., Xu, X., Kou, Y., Lv, Y., Li, Y., Liu, Y., Wang, Y., Zhang, Y., Huang, Y., Li, Y., Wu, Y., Liu, Y., Pan, Y., Zheng, Y., Hong, Y., Shi, Y., Feng, Y., Jiang, Z., Han, Z., Wu, Z.F., Liu, Z.: Wan: Open and advanced large-scale video generative models (2025), <https://arxiv.org/abs/2503.20314>
 40. Wang, B., Huang, C., Wen, J., Wang, W., Liu, Y., Xu, Y.: Federated weakly supervised video anomaly detection with multimodal prompt. In: Proceedings of the AAAI conference on artificial intelligence. vol. 39, pp. 21017–21025 (2025)
 41. Wang, X., Che, Z., Jiang, B., Xiao, N., Yang, K., Tang, J., Ye, J., Wang, J., Qi, Q.: Robust unsupervised video anomaly detection by multipath frame prediction. IEEE transactions on neural networks and learning systems **33**(6), 2301–2312 (2021)
 42. Wiedemer, T., Li, Y., Vicol, P., Gu, S.S., Matarese, N., Swersky, K., Kim, B., Jaini, P., Geirhos, R.: Video models are zero-shot learners and reasoners (2025), <https://arxiv.org/abs/2509.20328>
 43. Wu, P., Liu, J., Shi, Y., Sun, Y., Shao, F., Wu, Z., Yang, Z.: Not only look, but also listen: Learning multimodal violence detection under weak supervision. In: European conference on computer vision. pp. 322–339. Springer (2020)
 44. Wu, P., Zhou, X., Pang, G., Sun, Y., Liu, J., Wang, P., Zhang, Y.: Open-vocabulary video anomaly detection (2024), <https://arxiv.org/abs/2311.07042>
 45. Wu, P., Zhou, X., Pang, G., Zhou, L., Yan, Q., Wang, P., Zhang, Y.: Vadclip: Adapting vision-language models for weakly supervised video anomaly detection (2023), <https://arxiv.org/abs/2308.11681>

46. Yang, Y., Lee, K., Dariush, B., Cao, Y., Lo, S.Y.: Follow the rules: Reasoning for video anomaly detection with large language models (2024), <https://arxiv.org/abs/2407.10299>
47. Zanella, L., Menapace, W., Mancini, M., Wang, Y., Ricci, E.: Harnessing large language models for training-free video anomaly detection (2024), <https://arxiv.org/abs/2404.01014>
48. Zhang, H., Xu, X., Wang, X., Zuo, J., Huang, X., Gao, C., Zhang, S., Yu, L., Sang, N.: Holmes-vau: Towards long-term video anomaly understanding at any granularity (2025), <https://arxiv.org/abs/2412.06171>
49. Zhang, H., Xu, X., Wang, X., Zuo, J., Huang, X., Gao, C., Zhang, S., Yu, L., Sang, N.: Holmes-vau: Towards long-term video anomaly understanding at any granularity. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 13843–13853 (2025)
50. Zhang, M., Wang, J., Qi, Q., Sun, H., Zhuang, Z., Ren, P., Ma, R., Liao, J.: Multi-scale video anomaly detection by multi-grained spatio-temporal representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17385–17394 (2024)
51. Zhang, Y., Zhou, D., Chen, S., Gao, S., Ma, Y.: Single-image crowd counting via multi-column convolutional neural network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 589–597 (2016)
52. Zhang, Y., Li, B., Liu, h., Lee, Y.j., Gui, L., Fu, D., Feng, J., Liu, Z., Li, C.: Llava-next: A strong zero-shot video understanding model (April 2024), <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>
53. Zhou, H., Yu, J., Yang, W.: Dual memory units with uncertainty regulation for weakly supervised video anomaly detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3769–3777 (2023)
54. Zhou, H., Yu, J., Yang, W.: Dual memory units with uncertainty regulation for weakly supervised video anomaly detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3769–3777 (2023)
55. Zhu, L., Wang, L., Raj, A., Gedeon, T., Chen, C.: Advancing video anomaly detection: A concise review and a new dataset. *Advances in Neural Information Processing Systems* **37**, 89943–89977 (2024)