

Novelty Assessment Report

Paper: SPIKE-RL: Video-LLMs meet Bayesian Surprise

PDF URL: <https://openreview.net/pdf?id=QLiXtWEAkq>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-01

Abstract

Real-world videos often show routine activities punctuated by memorable, surprising events. However, most Video-LLMs process videos by sampling frames uniformly, likely missing critical moments that define a video's narrative. We introduce SPIKE, an inference-time framework that quantifies Bayesian Surprise as the belief update triggered by new visual evidence in the video stream, identifying moments where new visual evidence conflicts with prior beliefs. SPIKE effectively localizes surprise in videos, correlated with humans on positive (FunQA) and negative (Oops!) surprise benchmarks. SPIKE-RL further improves on SPIKE's ability to detect surprise, leveraging GRPO to refine its belief hypotheses based on a reward signal from the video caption. SPIKE and SPIKE-RL guide query-agnostic surprise-weighted frame sampling, which allocates more frames to interesting moments in the video. With this strategy, we achieve consistent performance gains on five downstream benchmarks. By enabling Video-LLMs to track beliefs and register surprise, our work paves the way for more robust models that can revise their understanding in response to new information.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Bayesian surprise detection in videos**

A total of **50 papers** were analyzed and organized into a taxonomy with **15 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Bayesian Surprise Theory and Computational Frameworks**
- **Video Anomaly and Event Detection**
- **Multimodal and Spatiotemporal Bayesian Modeling**
- **Robotics and Autonomous Systems Applications**
- **Specialized Application Domains**

Complete Taxonomy Tree

- Bayesian surprise detection in videos Survey Taxonomy
- Bayesian Surprise Theory and Computational Frameworks
 - Core Surprise Models and Attention Mechanisms (3 papers)
 - [9] A principled approach to detecting surprising events in video (Laurent Itti, 2005) [View paper](#)
 - [14] Modeling emotions associated with novelty at variable uncertainty levels: A Bayesian approach (Hideyoshi Yanagisawa, 2019) [View paper](#)
 - [24] Visual saliency detection using information divergence (Weilong Hou, 2013) [View paper](#)
 - Surprise-Based Exploration and Curiosity (3 papers)
 - [31] Curiosity-Driven Exploration via Latent Bayesian Surprise (Mazzaglia, 2021) [View paper](#)
 - [34] Self-Supervised Exploration via Latent Bayesian Surprise. (Pietro Mazzaglia, 2021) [View paper](#)
 - [35] Perception-driven curiosity with bayesian surprise (Bernadette Bucher, 2019) [View paper](#)
 - Surprise in Cognitive and Neural Systems (6 papers)
 - [2] Brain network dynamics predict moments of surprise across contexts (Ziwei Zhang, 2025) [View paper](#)
 - [32] Learning to infer in recurrent biological networks (Benjamin, 2020) [View paper](#)
 - [37] Subjective Depth and Timescale Transformers: Learning Where and When to Compute (Frederico Wieser, 2025) [View paper](#)
 - [38] Finding Similar Objects and Active Inference for Surprise in Numenta Neocortex Model (Kawakami Hajime, 2025) [View paper](#)
 - [43] Human-inspired Episodic Memory for Infinite Context LLMs (Z. Fountas, 2024) [View paper](#)
 - [50] Bayesian Surprise Predicts Human Event Segmentation in Story Listening. (Manoj Kumar, 2023) [View paper](#)
- Video Anomaly and Event Detection
 - Bayesian Nonparametric and Topic Models for Anomaly Detection (5 papers)
 - [7] Anomaly detection in video with Bayesian nonparametrics (Isupova Olga, 2016) [View paper](#)
 - [12] Anomaly detection in unstructured environments using Bayesian nonparametric scene modeling (Yogesh Girdhar, 2015) [View paper](#)
 - [33] Bayesian Nonparametric Approaches to Abnormality Detection in Video Surveillance (Vu Nguyen, 2015) [View paper](#)
 - [36] Machine learning methods for behaviour analysis and anomaly detection in video (Olga Isupova, 2018) [View paper](#)
 - [47] Bayesian Nonparametric Submodular Video Partition for Robust Anomaly Detection (Hitesh Sapkota, 2022) [View paper](#)
 - Hierarchical and Gaussian Process-Based Anomaly Detection (4 papers)
 - [1] Video anomaly detection and localization using hierarchical feature representation and Gaussian process regression (Kai-wen Cheng, 2015) [View paper](#)
 - [10] Video event recognition and anomaly detection by combining gaussian process and hierarchical dirichlet process models (Yang, 2018) [View paper](#)

- [25] Anomaly Detection in Video Surveillance via Gaussian Process (Nannan Li, 2015) [View paper](#)
- [29] Online nonparametric Bayesian activity mining and analysis from surveillance video (Vahid Bastani, 2016) [View paper](#)
- Deep Generative Models for Anomaly Detection (3 papers)
- [3] Multilevel Anomaly Detection Through Variational Autoencoders and Bayesian Models for Self-Aware Embodied Agents (Giulia Slavic, 2022) [View paper](#)
- [8] Bayesian Feed Forward Neural Network-Based Efficient Anomaly Detection from Surveillance Videos. (M Murugesan, 2022) [View paper](#)
- [13] Sparsity-based naive bayes approach for anomaly detection in real surveillance videos (Kamoona Ammar Mansoor, 2019) [View paper](#)
- Surprise-Based Event and Saliency Detection ★ (8 papers)
- [0] SPIKE-RL: Video-LLMs meet Bayesian Surprise (Anon et al., 2026) [View paper](#)
- [18] An audio-visual human attention analysis approach to abrupt change detection in videos (Yanxiang Chen, 2015) [View paper](#)
- [26] Neuromorphic Bayesian Surprise for Far-Range Event Detection (Randolph Voorhies, 2012) [View paper](#)
- [27] Identifying surprising events in videos using bayesian topic models (A. Hendel, 2010) [View paper](#)
- [39] Identifying Surprising Events in Video Using Bayesian Topic Models (Avishai Hendel, 2010) [View paper](#)
- [40] Bayesian Surprise for Small and Sub-Pixel Moving Target Detection (Laurent Itti, 2025) [View paper](#)
- [42] The detection of unusual events in video based on Bayesian surprise model (Jinsheng Xie, 2010) [View paper](#)
- [49] Application of a bottom-up visual surprise model for event detection in dynamic natural scenes (Randolph Voorhies, 2010) [View paper](#)
- Behavior Profiling and Specific Event Detection (4 papers)
- [6] Contextual Anomaly Detection Based Video Surveillance System (S. A. Mahmood, 2021) [View paper](#)
- [11] Specific event detection for video surveillance using variational Bayesian inference (Roberto Leyva, 2024) [View paper](#)
- [21] Video behavior profiling for anomaly detection (Tao Xiang, 2008) [View paper](#)
- [44] Fall Detection in Hospital Rooms with Probabilistic Graphical Models (Jacob González de Frutos, 2024) [View paper](#)
- Multimodal and Spatiotemporal Bayesian Modeling
 - Dynamic Bayesian Networks for Video and Sensor Fusion (4 papers)
 - [5] Simultaneous Localization and Anomaly Detection from First-Person Video Data through a Coupled Dynamic Bayesian Network Model (Giulia Slavic, 2022) [View paper](#)
 - [15] A Kalman Variational Autoencoder Model Assisted by Odometric Clustering for Video Frame Prediction and Anomaly Detection (Giulia Slavic, 2022) [View paper](#)
 - [30] Dynamic Bayesian Network-Based Anomaly Detection for In-Process Visual Inspection of Laser Surface Heat Treatment (A. Ogbechie, 2016) [View paper](#)
 - [48] Using dynamic Bayesian network for scene modeling and anomaly detection (Imran N. Junejo, 2010) [View paper](#)
 - Trajectory and Traffic Anomaly Detection (2 papers)
 - [16] Spatiotemporal multi-feature fusion vehicle trajectory anomaly detection for intelligent transportation: An improved method combining autoencoders and dynamic Bayesian networks. (Mingqi Qiu, 2025) [View paper](#)
 - [17] A Bayesian forecasting and anomaly detection framework for vehicular monitoring networks (Maria Scalabrin, 2017) [View paper](#)
- Robotics and Autonomous Systems Applications
 - Surprise-Based Landmark Detection and Localization (3 papers)
 - [41] Automatic Landmark Detection for Topological Mapping Using Bayesian Surprise (Ananth Ranganathan, 2008) [View paper](#)
 - [45] Bayesian surprise and landmark detection (Ananth Ranganathan, 2009) [View paper](#)
 - [46] A probabilistic appearance representation and its application to surprise detection in cognitive robots (Werner Maier, 2010) [View paper](#)
 - Anomaly Detection for Autonomous Vehicles (1 papers)
 - [28] Anomaly Detection for Autonomous Guided Vehicles using Bayesian Surprise (Ozan Aytal, 2020) [View paper](#)
- Specialized Application Domains
 - Human Behavior and Developmental Assessment (2 papers)
 - [19] Surprise! predicting infant visual attention in a socially assistive robot contingent learning paradigm (Lauren Klein, 2019) [View paper](#)
 - [23] Computer vision to automatically assess infant neuromotor risk (Claire Chambers, 2019) [View paper](#)
 - Discrete Sequence and Non-Video Anomaly Detection (2 papers)
 - [4] Cognitive Based Detection of Anomalous Sequences Using Bayesian Surprise (Ken McGarry, 2025) [View paper](#)
 - [22] Unsupervised detection of interesting discrete sequences using Bayesian surprise (Ken McGarry, 2022) [View paper](#)
 - Recommendation Systems and Content Analysis (1 papers)
 - [20] A Deep Bayesian Tensor-Based System for Video Recommendation (Wei Lu, 2018) [View paper](#)

Narrative

Core task: Bayesian surprise detection in videos. The field centers on identifying unexpected or anomalous events in video streams by quantifying deviations from learned probabilistic models of normal behavior. The taxonomy reveals five main branches that reflect different emphases: Bayesian Surprise Theory and Computational Frameworks develops the mathematical foundations and inference algorithms (e.g., Hierarchical Gaussian Process[1], Variational Bayesian Inference[11]); Video Anomaly and Event Detection applies these principles to surveillance and event recognition (e.g., Contextual Video Surveillance[6], Unusual Events Detection[42]); Multimodal and Spatiotemporal Bayesian Modeling extends surprise measures across sensory modalities and temporal scales (e.g., Audio-Visual Attention Analysis[18], Kalman Variational Autoencoder[15]); Robotics and Autonomous Systems Applications leverages surprise for navigation and decision-making (e.g., Autonomous Vehicle Surprise[28], Landmark Bayesian Surprise[41]); and Specialized Application Domains targets niche settings such as healthcare monitoring (Fall Detection Hospitals[44]) or cognitive modeling (Infant Visual Attention[19]). Together, these branches illustrate a progression from theoretical constructs to diverse real-world deployments.

A particularly active line of work explores how surprise-based saliency and attention mechanisms can guide both bottom-up perceptual processing (Bottom-Up Visual Surprise[49]) and top-down event segmentation (Bayesian Topic Events[39]). Trade-offs emerge between computational efficiency—favoring lightweight neuromorphic implementations (Neuromorphic Bayesian Surprise[26])—and representational richness in deep generative models (Multilevel Variational Autoencoders[3]). SPIKE-RL[0] sits within the Surprise-Based Event and Saliency Detection cluster, emphasizing reinforcement learning integration for dynamic video analysis. Its approach contrasts with purely unsupervised anomaly detectors like Simultaneous Localization Anomaly[5], which focus on spatial consistency, and with classical information-theoretic methods such as Information Divergence Saliency[24], which lack adaptive learning. By combining

Bayesian surprise with RL, SPIKE-RL[0] bridges perceptual novelty detection and goal-directed behavior, positioning itself at the intersection of event detection and autonomous decision-making.

Related Works in Same Category

The following 7 **sibling papers** share the same taxonomy leaf node with the original paper:

1. An audio-visual human attention analysis approach to abrupt change detection in videos

Authors: Yanxiang Chen, Minglong Song, Lixia Xue, Xiaoxue Chen, Meng Wang, et al. (6 authors total) | **Year/Venue:** 2015 | **URL:** [View paper](#)

Abstract

We detect abrupt changes in videos based on Bayesian surprise theory, which considers both visual and audio modalities. Specifically, after generating surprise visual series in videos. The

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify salient moments in video streams. The candidate paper focuses on audiovisual human attention analysis for abrupt change detection by combining visual and audio modalities with Bayesian surprise theory. In contrast, the original paper (SPIKE-RL) integrates Video-LLMs with Bayesian surprise for belief tracking and uses reinforcement learning to optimize hypothesis generation, targeting surprise-weighted frame sampling for downstream video understanding tasks rather than abrupt change detection.

2. Neuromorphic Bayesian Surprise for Far-Range Event Detection

Authors: Randolph Voorhies, Lior Elazary, Randolph C. Voorhies, L. Itti, Laurent Itti | **Year/Venue:** 2012 • 2012 IEEE Ninth International Conference on Advanced Video and Signal-Based Surveillance | **URL:** [View paper](#)

Abstract

N/A

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify salient or unexpected events in video streams. The original paper (SPIKE-RL) focuses on integrating Video-LLMs with Bayesian surprise for frame sampling and downstream video understanding tasks, while the candidate paper addresses far-range event detection in surveillance footage using neuromorphic Bayesian surprise computed over multiple spatial and temporal scales with low-level visual features. The key difference is that SPIKE-RL operates at a semantic level with language models and belief hypotheses, whereas the candidate paper works at a low-level feature detection level for small, distant targets in high-resolution surveillance scenarios.

3. Identifying surprising events in videos using bayesian topic models

Authors: A. Hendel, D. Weinshall, Avishai Hendel, Shmuel Peleg, Daphna Weinshall | **Year/Venue:** 2010 | **URL:** [View paper](#)

Abstract

We video events, which are the events a human user is usually looking for. To this end we employ the notion of Bayesian surprise. The results demonstrate our ability to detect atypical events.

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify unusual events in videos. They overlap in applying Bayesian surprise (KL divergence between prior and posterior beliefs) to video analysis, but differ fundamentally in their approach: the original paper (SPIKE-RL) uses Video-LLMs with textual belief hypotheses and reinforcement learning for surprise-weighted frame sampling across diverse video understanding tasks, while the candidate paper uses LDA topic models over trajectory features to detect atypical activities in surveillance videos of street intersections.

4. Identifying Surprising Events in Video Using Bayesian Topic Models

Authors: Avishai Hendel, Daphna Weinshall, Bezalel Peleg | **Year/Venue:** 2010 • Detection and Identification of Rare Audiovisual Cues | **URL:** [View paper](#)

Abstract

N/A

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify interesting events in videos. The candidate paper focuses on using Bayesian topic models for surprise detection, while the original paper (SPIKE-RL) introduces a Video-LLM-based framework that tracks beliefs as textual hypotheses and uses reinforcement learning to optimize surprise detection. The key difference is that SPIKE-RL leverages large language models with explicit belief tracking and RL-based refinement, whereas the candidate appears to use traditional topic modeling approaches for surprise identification.

5. Bayesian Surprise for Small and Sub-Pixel Moving Target Detection

Authors: Laurent Itti, Daben Liu, L. Itti, Corinne Teeter, Srideep Musuvathy, et al. (6 authors total) | **Year/Venue:** 2025 | **URL:** [View paper](#)

Abstract

We Bayesian surprise. It computes how an artificial observer's beliefs about the contents of a video. We hypothesize that unpredictable but repetitive events, such as foliage fluttering in the

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify salient events in video streams. While the original paper (SPIKE-RL) focuses on using Bayesian surprise with Video-LLMs to detect semantically surprising narrative moments (like unexpected falls or humorous events) and improve video understanding tasks, the candidate paper applies Bayesian surprise specifically to the technical problem of detecting small and sub-pixel moving targets in video, addressing challenges like foliage movement and low signal-to-noise scenarios. The key difference lies in their application domains: semantic event understanding versus low-level motion detection.

6. The detection of unusual events in video based on Bayesian surprise model

Authors: Jinsheng Xie, Li Guo, Yunbi Chen, Jin-Sheng Xie, Long Zhao, et al. (6 authors total) | **Year/Venue:** 2010 • International Conference on Information Science and Engineering | **URL:** [View paper](#)

Abstract

N/A

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify unusual or salient events in video streams. The candidate paper appears to focus on detecting unusual events using a Bayesian surprise model, overlapping with the original paper's core approach of quantifying Bayesian surprise for video analysis. However, the original paper (SPIKE-RL) specifically integrates Video-LLMs with belief tracking, reinforcement learning optimization, and surprise-weighted frame sampling for downstream tasks, while the candidate paper's specific technical approach and application scope remain unclear from the limited abstract information provided.

7. Application of a bottom-up visual surprise model for event detection in dynamic natural scenes

Authors: Randolph Voorhies, Lior Elazary, R. Voorhies, L. Itti | **Year/Venue:** 2010 | **URL:** [View paper](#)

Abstract

We present an application of a neuromorphic visual attention model to the field of large-scale video surveillance and show that it outperforms a state-of-the-art method at the task of event detection. Our work extends Itti and Baldi's Surprise framework as described by "A Principled Approach to Detecting Surprising Events in Video" in CVPR 2005. The Surprise framework is a biologically plausible and validated model of primate visual attention which uses a new Bayesian model of information to...

Relationship Analysis

Both papers belong to the Surprise-Based Event and Saliency Detection category, using Bayesian surprise metrics to identify unexpected events in video streams. The candidate paper applies Itti and Baldi's Bayesian surprise framework to large-scale video surveillance for event detection in outdoor environments, focusing on detecting pedestrians and vehicles using biologically-inspired visual attention models. In contrast, the original paper (SPIKE-RL) integrates Bayesian surprise with modern Video-LLMs to track beliefs through textual hypotheses and uses reinforcement learning to optimize surprise detection, targeting diverse video understanding tasks rather than surveillance-specific applications.

Contributions Analysis

Overall novelty summary. The paper introduces SPIKE, a framework that quantifies Bayesian surprise in video streams to identify moments where new visual evidence conflicts with prior beliefs, and SPIKE-RL, which refines belief hypotheses using reinforcement learning. Within the taxonomy, this work resides in the 'Surprise-Based Event and Saliency Detection' leaf under 'Video Anomaly and Event Detection', alongside seven sibling papers. This leaf represents a moderately populated research direction focused on using explicit surprise metrics for event identification, distinguishing it from broader anomaly detection approaches that rely on topic models or Gaussian processes without surprise-based formulations.

The taxonomy reveals that neighboring leaves include 'Bayesian Nonparametric and Topic Models for Anomaly Detection' (five papers using Dirichlet processes) and 'Deep Generative Models for Anomaly Detection' (three papers combining VAEs with Bayesian methods). The paper's approach diverges from these by emphasizing inference-time surprise computation rather than offline model training, and by integrating reinforcement learning for belief optimization. Its connection to 'Core Surprise Models and Attention Mechanisms' (three foundational papers) suggests it builds on established surprise theory while extending it to modern Video-LLM architectures, bridging classical Bayesian frameworks with contemporary deep learning systems.

Among twenty candidates examined across three contributions, the SPIKE framework shows one refutable candidate out of ten examined, indicating some prior work on Bayesian surprise quantification exists within the limited search scope. SPIKE-RL, however, encountered zero refutable candidates among ten examined, suggesting its reinforcement learning approach for belief refinement may represent a less-explored direction. The surprise-weighted frame sampling strategy was not evaluated against prior work (zero candidates examined), leaving its novelty assessment incomplete. These statistics reflect a focused semantic search rather than exhaustive coverage, so additional related work may exist beyond the top-twenty matches.

Based on the limited search scope of twenty semantically similar papers, the work appears to occupy a moderately novel position. The core surprise detection mechanism has some precedent, but the integration with reinforcement learning and application to Video-LLM frame sampling shows less overlap with examined candidates. The taxonomy structure indicates this is an active but not overcrowded research area, with the sibling leaf containing eight papers total. A more comprehensive literature review would be needed to assess whether the specific combination of Bayesian surprise, RL-based belief optimization, and LLM-guided sampling has been explored elsewhere.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: SPIKE inference-time framework for Bayesian Surprise quantification

Description: SPIKE is a framework that represents a Video-LLM's beliefs as explicit probability distributions over textual hypotheses and measures surprise as the KL divergence between prior and posterior beliefs when new frames are observed. This enables the model to identify moments where visual evidence conflicts with expectations.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Electroencephalographic correlates of temporal Bayesian belief updating and surprise

URL: [View paper](#)

Brief Assessment

Electroencephalographic Bayesian Updating[67] focuses on neural correlates of temporal prediction updating in EEG data, not on video understanding or Video-LLM belief tracking frameworks.

2. Rejecting outliers: Surprising changes do not always improve belief updating.

URL: [View paper](#)

Brief Assessment

Rejecting Outliers[64] studies human belief updating in a Plinko game task, measuring how surprising changes affect learning. It does not propose a Video-LLM framework or use KL divergence between textual hypothesis distributions for video understanding.

3. Uncertainty and persistence: A Bayesian update semantics for probabilistic expressions

URL: [View paper](#)

Brief Assessment

Uncertainty and Persistence[66] focuses on linguistic semantics and Bayesian updates for probabilistic expressions in natural language, not on video understanding or surprise quantification in visual streams.

4. AutoDiscovery: Open-ended Scientific Discovery via Bayesian Surprise

URL: [View paper](#)

Brief Assessment

AutoDiscovery[62] applies Bayesian surprise to scientific hypothesis discovery across datasets, not to video understanding or belief tracking in video-LLMs. The domains and applications are fundamentally different.

5. Modeling emotions associated with novelty at variable uncertainty levels: A Bayesian approach

URL: [View paper](#)

Brief Assessment

Modeling Novelty Emotions[14] focuses on modeling emotional responses (arousal and valence) to novelty in human perception experiments using percussion instruments, not on video understanding or LLM belief tracking systems.

6. Bayesian surprise attracts human attention

URL: [View paper](#)

Prior Art Analysis

Bayesian Surprise Attention[63] demonstrates that the core concept of quantifying Bayesian surprise as the difference between posterior and prior beliefs was established prior to the ORIGINAL paper. The candidate explicitly describes 'a formal Bayesian definition of surprise' that 'quantifies how data affects a natural or artificial observer, by measuring the difference between posterior and prior beliefs of the observer.' This is the exact same mathematical formulation that SPIKE claims as novel. The candidate further validates this approach empirically by showing that 'subjects are strongly attracted towards surprising locations' in video content, demonstrating that Bayesian surprise computed via belief updates successfully identifies salient moments in videos - the same application domain as SPIKE.

Evidence

Evidence 1 - **Rationale:** Both papers define Bayesian surprise identically: as measuring the difference/divergence between posterior and prior beliefs. The candidate establishes this as 'the only consistent formulation under minimal axiomatic assumptions,' indicating it was already a formal, established framework before SPIKE. - **Original:** we introduce spike, an inference-time framework that represents a model's beliefs as explicit probability distributions over human-interpretable textual hypotheses, and quantifies bayesian surprise as the divergence between prior and posterior beliefs - **Candidate:** we describe a formal bayesian definition of surprise that is the only consistent formulation under minimal axiomatic assumptions. surprise quantifies how data affects a natural or artificial observer, by measuring the difference between posterior and prior beliefs of the observer.

Evidence 2 - **Rationale:** The ORIGINAL paper explicitly cites the candidate work (Itti & Baldi 2005) as the source of their Bayesian surprise formalization. This citation directly acknowledges that the mathematical framework for quantifying surprise via KL divergence between beliefs was not novel to SPIKE but was adopted from prior work. - **Original:** following the bayesian formalization of surprise by itti & baldi (2005), we quantify our surprise score to be the information gain induced by ot, as the kullback-leibler (kl) divergence between posterior and prior beliefs over hypotheses - **Candidate:** we describe a formal bayesian definition of surprise that is the only consistent formulation under minimal axiomatic assumptions. surprise quantifies how data affects a natural or artificial observer, by measuring the difference between posterior and prior beliefs of the observer.

7. Neural signals encoding shifts in beliefs

URL: [View paper](#)

Brief Assessment

Neural Belief Shifts[65] focuses on neural encoding of belief shifts in the brain, not on computational frameworks for video-LLMs. The candidate studies biological neural signals rather than machine learning inference methods.

8. Brain network dynamics predict moments of surprise across contexts

URL: [View paper](#)

Brief Assessment

Brain Network Dynamics[2] applies Bayesian surprise to brain imaging data for predicting mental states, not to video-LLM belief tracking or frame sampling in video understanding tasks.

9. Hierarchical surprise signals in naturalistic violation of expectations

URL: [View paper](#)

Brief Assessment

Hierarchical Surprise Signals[61] focuses on fMRI neural responses to magic tricks in human brains, not computational frameworks for Video-LLMs. The candidate studies biological surprise processing in cognitive neuroscience, while the original develops a computational belief-tracking system for video understanding models.

10. Bayesian Surprise for Small and Sub-Pixel Moving Target Detection

URL: [View paper](#)

Brief Assessment

Small Target Detection[40] focuses on detecting small and sub-pixel moving targets in videos using Bayesian surprise for motion detection, not on belief tracking or hypothesis generation for video understanding tasks.

Contribution 2: SPIKE-RL reinforcement learning method for belief optimization

Description: SPIKE-RL uses Group Relative Policy Optimization (GRPO) to train the hypothesis generator by propagating rewards from final caption quality back to intermediate belief hypotheses. This improves both the diversity of generated beliefs and the accuracy of surprise localization beyond the inference-time scorer alone.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Avatar: Reinforcement learning to see, hear, and reason over video

URL: [View paper](#)

Brief Assessment

Avatar[56] uses GRPO for multimodal video reasoning with temporal advantage shaping, not for belief hypothesis optimization in Bayesian surprise frameworks. The technical focus and application domain differ fundamentally from SPIKE-RL's belief tracking approach.

2. Thinking with videos: Multimodal tool-augmented reinforcement learning for long video reasoning

URL: [View paper](#)

Brief Assessment

Thinking with Videos[58] uses GRPO for tool-augmented video reasoning with visual toolbox integration, not for belief hypothesis optimization in Bayesian surprise detection. The technical focus and application domains differ fundamentally.

3. The social machine: artificial intelligence (AI) approaches to theory of mind

URL: [View paper](#)

Brief Assessment

Social Machine AI[59] appears to focus on theory of mind and belief reasoning in social contexts, while the original paper's SPIKE-RL specifically addresses video understanding through Bayesian surprise detection and frame sampling optimization. The candidate's limited context does not demonstrate prior work on using GRPO for belief hypothesis optimization in video models.

4. Videohallu: Evaluating and mitigating multi-modal hallucinations on synthetic video understanding

URL: [View paper](#)

Brief Assessment

Videohallu[52] focuses on evaluating and mitigating hallucinations in synthetic video understanding using GRPO for post-training, not on belief hypothesis optimization for surprise detection in video streams.

5. Spatialladder: Progressive training for spatial reasoning in vision-language models

URL: [View paper](#)

Brief Assessment

Spatialladder[53] focuses on progressive training for spatial reasoning in vision-language models through three stages (perception, understanding, reasoning), not on belief hypothesis optimization in video models. The candidate uses GRPO for spatial task optimization, while the original applies it to belief trajectory refinement for surprise detection.

6. Exploring the effect of reinforcement learning on video understanding: Insights from seed-bench-r1

URL: [View paper](#)

Brief Assessment

Seed-Bench Insights[54] applies GRPO to video understanding tasks but does not address belief hypothesis generation or Bayesian surprise quantification. The candidate focuses on action planning in egocentric videos, not belief tracking mechanisms.

7. Video-Thinker: Sparking "Thinking with Videos" via Reinforcement Learning

URL: [View paper](#)

Brief Assessment

Video-Thinker[57] applies GRPO to optimize video reasoning with grounding/captioning capabilities, not belief hypothesis generation for surprise detection. The technical focus and application domain differ fundamentally from SPIKE-RL's Bayesian surprise framework.

8. RLZero: Direct Policy Inference from Language Without In-Domain Supervision

URL: [View paper](#)

Brief Assessment

RLZero[60] focuses on zero-shot policy inference from language instructions using pretrained RL agents without task-specific supervision, not on belief hypothesis optimization for video understanding or surprise detection.

9. ReAgent-V: A Reward-Driven Multi-Agent Framework for Video Understanding

URL: [View paper](#)

Brief Assessment

ReAgent-V[51] uses GRPO for video understanding tasks with multi-agent reflection mechanisms, not for belief hypothesis optimization in Bayesian surprise detection frameworks.

10. Video-rts: Rethinking reinforcement learning and test-time scaling for efficient and enhanced video reasoning

URL: [View paper](#)

Brief Assessment

Video-RTS[55] focuses on outcome-supervised RL (GRPO) for video question-answering without belief tracking or hypothesis generation. The candidate does not address belief optimization or surprise localization mechanisms.

Contribution 3: Surprise-weighted frame sampling strategy for Video-LLMs

Description: The authors propose replacing uniform frame sampling in Video-LLMs with a surprise-weighted sampling strategy that allocates the frame budget proportionally to computed surprise scores. This query-agnostic approach consistently improves performance on downstream video understanding tasks.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] SPIKE-RL: Video-LLMs meet Bayesian Surprise [View paper](#)
- [1] Video anomaly detection and localization using hierarchical feature representation and Gaussian process regression [View paper](#)
- [2] Brain network dynamics predict moments of surprise across contexts [View paper](#)
- [3] Multilevel Anomaly Detection Through Variational Autoencoders and Bayesian Models for Self-Aware Embodied Agents [View paper](#)
- [4] Cognitive Based Detection of Anomalous Sequences Using Bayesian Surprise [View paper](#)
- [5] Simultaneous Localization and Anomaly Detection from First-Person Video Data through a Coupled Dynamic Bayesian Network Model [View paper](#)

- [6] Contextual Anomaly Detection Based Video Surveillance System [View paper](#)
- [7] Anomaly detection in video with Bayesian nonparametrics [View paper](#)
- [8] Bayesian Feed Forward Neural Network-Based Efficient Anomaly Detection from Surveillance Videos. [View paper](#)
- [9] A principled approach to detecting surprising events in video [View paper](#)
- [10] Video event recognition and anomaly detection by combining gaussian process and hierarchical dirichlet process models [View paper](#)
- [11] Specific event detection for video surveillance using variational Bayesian inference [View paper](#)
- [12] Anomaly detection in unstructured environments using Bayesian nonparametric scene modeling [View paper](#)
- [13] Sparsity-based naive bayes approach for anomaly detection in real surveillance videos [View paper](#)
- [14] Modeling emotions associated with novelty at variable uncertainty levels: A Bayesian approach [View paper](#)
- [15] A Kalman Variational Autoencoder Model Assisted by Odometric Clustering for Video Frame Prediction and Anomaly Detection [View paper](#)
- [16] Spatiotemporal multi-feature fusion vehicle trajectory anomaly detection for intelligent transportation: An improved method combining autoencoders and dynamic Bayesian networks. [View paper](#)
- [17] A Bayesian forecasting and anomaly detection framework for vehicular monitoring networks [View paper](#)
- [18] An audio-visual human attention analysis approach to abrupt change detection in videos [View paper](#)
- [19] Surprise! predicting infant visual attention in a socially assistive robot contingent learning paradigm [View paper](#)
- [20] A Deep Bayesian Tensor-Based System for Video Recommendation [View paper](#)
- [21] Video behavior profiling for anomaly detection [View paper](#)
- [22] Unsupervised detection of interesting discrete sequences using Bayesian surprise [View paper](#)
- [23] Computer vision to automatically assess infant neuromotor risk [View paper](#)
- [24] Visual saliency detection using information divergence [View paper](#)
- [25] Anomaly Detection in Video Surveillance via Gaussian Process [View paper](#)
- [26] Neuromorphic Bayesian Surprise for Far-Range Event Detection [View paper](#)
- [27] Identifying surprising events in videos using bayesian topic models [View paper](#)
- [28] Anomaly Detection for Autonomous Guided Vehicles using Bayesian Surprise [View paper](#)
- [29] Online nonparametric bayesian activity mining and analysis from surveillance video [View paper](#)
- [30] Dynamic Bayesian Network-Based Anomaly Detection for In-Process Visual Inspection of Laser Surface Heat Treatment [View paper](#)
- [31] Curiosity-Driven Exploration via Latent Bayesian Surprise [View paper](#)
- [32] Learning to infer in recurrent biological networks [View paper](#)
- [33] Bayesian Nonparametric Approaches to Abnormality Detection in Video Surveillance [View paper](#)
- [34] Self-Supervised Exploration via Latent Bayesian Surprise. [View paper](#)
- [35] Perception-driven curiosity with bayesian surprise [View paper](#)
- [36] Machine learning methods for behaviour analysis and anomaly detection in video [View paper](#)
- [37] Subjective Depth and Timescale Transformers: Learning Where and When to Compute [View paper](#)
- [38] Finding Similar Objects and Active Inference for Surprise in Numenta Neocortex Model [View paper](#)
- [39] Identifying Surprising Events in Video Using Bayesian Topic Models [View paper](#)
- [40] Bayesian Surprise for Small and Sub-Pixel Moving Target Detection [View paper](#)
- [41] Automatic Landmark Detection for Topological Mapping Using Bayesian Surprise [View paper](#)
- [42] The detection of unusual events in video based on Bayesian surprise model [View paper](#)
- [43] Human-inspired Episodic Memory for Infinite Context LLMs [View paper](#)
- [44] Fall Detection in Hospital Rooms with Probabilistic Graphical Models [View paper](#)
- [45] Bayesian surprise and landmark detection [View paper](#)
- [46] A probabilistic appearance representation and its application to surprise detection in cognitive robots [View paper](#)
- [47] Bayesian Nonparametric Submodular Video Partition for Robust Anomaly Detection [View paper](#)
- [48] Using dynamic Bayesian network for scene modeling and anomaly detection [View paper](#)
- [49] Application of a bottom-up visual surprise model for event detection in dynamic natural scenes [View paper](#)
- [50] Bayesian Surprise Predicts Human Event Segmentation in Story Listening. [View paper](#)
- [51] ReAgent-V: A Reward-Driven Multi-Agent Framework for Video Understanding [View paper](#)
- [52] Videohallu: Evaluating and mitigating multi-modal hallucinations on synthetic video understanding [View paper](#)
- [53] Spatialladder: Progressive training for spatial reasoning in vision-language models [View paper](#)
- [54] Exploring the effect of reinforcement learning on video understanding: Insights from seed-bench-r1 [View paper](#)
- [55] Video-rts: Rethinking reinforcement learning and test-time scaling for efficient and enhanced video reasoning [View paper](#)
- [56] Avatar: Reinforcement learning to see, hear, and reason over video [View paper](#)
- [57] Video-Thinker: Sparking" Thinking with Videos" via Reinforcement Learning [View paper](#)
- [58] Thinking with videos: Multimodal tool-augmented reinforcement learning for long video reasoning [View paper](#)
- [59] The social machine: artificial intelligence (AI) approaches to theory of mind [View paper](#)
- [60] RLZero: Direct Policy Inference from Language Without In-Domain Supervision [View paper](#)
- [61] Hierarchical surprise signals in naturalistic violation of expectations [View paper](#)
- [62] AutoDiscovery: Open-ended Scientific Discovery via Bayesian Surprise [View paper](#)
- [63] Bayesian surprise attracts human attention [View paper](#)
- [64] Rejecting outliers: Surprising changes do not always improve belief updating. [View paper](#)
- [65] Neural signals encoding shifts in beliefs [View paper](#)
- [66] Uncertainty and persistence: A Bayesian update semantics for probabilistic expressions [View paper](#)
- [67] Electroencephalographic correlates of temporal Bayesian belief updating and surprise [View paper](#)