

NAACA: Training-Free NeuroAuditory Attentive Cognitive Architecture with Oscillatory Working Memory for Saliency-Driven Attention Gating

Anonymous Authors¹

Abstract

Audio provides critical situational cues, yet current Audio Language Models (ALMs) face an attention bottleneck in long-form recordings where dominant background patterns can dilute rare, salient events. We introduce NAACA, a training-free NeuroAuditory Attentive Cognitive Architecture that reframes attention allocation as an auditory saliency filtering problem. At its core is OWM, a neuro-inspired Oscillatory Working Memory that maintains stable attractor-like states and triggers higher-cognition ALM processing only when adaptive energy fluctuations signal perceptual saliency, triggering higher-level reasoning. On XD-Violence, NAACA improves AudioQwen’s average precision (AP) from 53.50% to 70.60% while reducing unnecessary ALM invocations. Furthermore, qualitative case studies on the Urban Soundscapes of the World (USoW) dataset show that OWM captures novel events and subcategory shifts while remaining robust to transient pauses and ambient urban noise.

1. Introduction

Audio provides critical situational cues when vision is degraded or unavailable, enabling detection of abnormal events in public spaces such as distress calls, fights, or sudden crowd escalation, while also supporting acoustic monitoring in natural environments, where shifts in soundscapes can indicate biodiversity changes. Recent Audio Language Models (ALMs) have advanced audio understanding beyond narrow classification, but face a critical pre-attentive bottleneck in real-time long-form analysis. Unlike biological hearing, which instinctively filters stable background soundscapes to prioritize salient stimuli, transformer-based ALMs allocate attention inefficiently over long contexts. When

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

applied to long audio streams, ALM based inference can miss rare but critical events that occur late in the recording, because model focus is dominated by earlier or persistent background patterns, an effect we refer to as attention dilution. Fig. 1 illustrates this failure mode on a representative example. While exhaustive processing of short segments could recover these scene subsets, the computational cost is prohibitive for continuous monitoring. Consequently, a fundamental trade-off arises between perceptual saliency recall and computational efficiency.

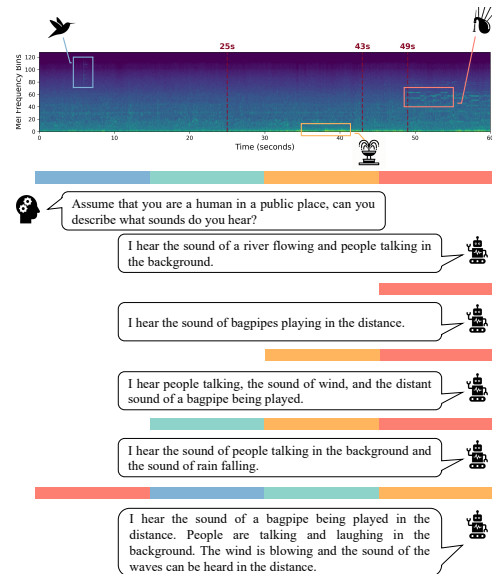


Figure 1. ALM attention failure and context limitations in long-form audio. **Top:** Mel-spectrogram of sample R0056 (USoW) with three salient scenes: birdsong (blue), increased fountain noise (yellow), and bagpipe onset (red). **Middle:** Direct inference on the full 60s clip (partitioned into 15s segments) omits the terminal bagpipe event, illustrating context-length bottleneck. **Bottom:** Varying context lengths and ordering. Trailing windows (15s–45s) recover different scene subsets; moving the final 15s to the front (reordered 60s) recovers the bagpipe, confirming temporal-attention decay in the Audio-LLM.

We address this trade-off by drawing inspiration from cognitive neuroscience, where selective attention mechanisms prioritize salient stimuli while filtering stable backgrounds. In auditory neuroscience, saliency refers to unexpected or behaviorally relevant stimuli (Corbetta & Shulman, 2002; Kayser et al., 2005). Crucially, saliency is not loudness: it

corresponds to context changes requiring internal updates (Friston, 2009), while stable backgrounds should be maintained without consuming attention (Winkler et al., 2009). Salient events are thus those warranting increased ALM focus.

Building on this saliency framework, we reformulate long-form processing as Saliency-Driven Attention Gating, an online segment selection policy that routes salient audio regions to the ALM. Rather than processing every window, we detect when the auditory pattern shifts in a salient manner and selectively invoke ALM reasoning only for those segments. This approach is critical for deployment, where offline training is prohibitive and data streams are unlabeled and non-stationary (Wan et al., 2024; Chan et al., 2025). Existing approaches, statistical detectors (Rabanser et al., 2019; Chan et al., 2025) and representation-based methods (Wan et al., 2024), require long-term historical data and substantial overhead.

To avoid the long-term historical data and substantial overhead required by existing methods, we draw inspiration from biological working memory. Working memory is supported by attractor-like neural states enabling stable maintenance while remaining sensitive to deviations requiring updates (Brennan & Proekt, 2023). Oscillatory dynamics link to attention and memory control, suggesting saliency emerges from state transitions rather than learned classification (Lundqvist et al., 2018).

Concretely, we propose **OWM (Oscillatory Working Memory)**, a bio-inspired module integrating frequency-selective oscillatory inputs from pretrained audio representations, maintaining stable internal states as attractor-like memory items, and identifying saliency by comparing energy fluctuations to an adaptive threshold. Building on OWM, we propose **NAACA**, a NeuroAuditory Attentive Architecture performing training-free online saliency gating for ALMs. As shown in Fig. 2, we process audio in sliding windows and convert each window into auditory object probability sequences, which drive frequency selective oscillatory inputs on the OWM grids. OWM maintains stable representations and triggers attentional focus when energy exceeds an adaptive threshold; detected segments are forwarded to an ALM for semantic interpretation. This enables operation without offline training while improving long-context reasoning by prioritizing salient transitions. On XD-Violence, our gating improves AudioQwen for 17.10 percentage points in AP while substantially reducing ALM invocations.

Our contributions are:

- **Neuroscience-Inspired Gating:** We formalize the ALM accuracy–cost trade-off and address it via saliency-driven gating.

- **Oscillatory Working Memory:** A bio-inspired module with optimized wave dynamics for training-free, adaptive saliency filtering.
- **Performance:** 70.60% AP with near 40% cost reduction and interpretable, neurally-plausible dynamics.

2. Methods

2.1. NeuroAuditory Attentive Cognitive Architecture (NAACA)

Our framework follows a multi-stage processing pipeline (Fig. 2 and Algorithm C.1). Incoming audio streams are segmented into short, overlapping windows \mathbf{x}_t , which are encoded into auditory object category probability vectors $\mathbf{p}_t = \text{Enc}(\mathbf{x}_t)$ by a pretrained encoder $\text{Enc}(\cdot)$.

These probabilities are transformed into oscillatory drive signals through a predefined modulator. Specifically, each probability dimension is assigned a unique carrier frequency, represented as a sine wave, while the corresponding probability value modulates its amplitude. As illustrated in Fig. 2, $\text{Enc}(\cdot)$ produces sound-category probability vectors across sliding windows, which are then mapped to their respective sinusoidal drive signals at distinct frequencies. Formally, the modulated source term for auditory object category i is

$$S_i(x, t) = a_i(t) \sin(\omega_i t) \mathbf{1}_{\Omega_i}(x), \tag{1}$$

$$\omega_i = 2\pi f_i, \quad a_i(t) \in [0, 1],$$

where $a_i(t)$ denotes the instantaneous amplitude given by the encoder probability for category i at time t , f_i is the carrier frequency assigned to category i (assignment details in Appendix B), and $\omega_i = 2\pi f_i$. $\Omega_i \subset \{1, \dots, G\} \times \{1, \dots, G\}$ is the spatial parcel on a $G \times G$ OWM lattice to which the i -th probability dimension $\mathbf{p}_t[i]$ injects an oscillatory drive, and $\mathbf{1}_{\Omega_i}(x)$ restricts the forcing to that region. Importantly, Eq. (1) only defines the *input*; wave-like propagation is produced by the OWM recurrent dynamics introduced next, where spatial coupling terms (via ∇p and $\nabla \cdot v$) allow activity induced in Ω_i to spread and interact across the grid under damping and wave propagation speed parameters. In this driven–damped setting, Ω_i serves as a consistent locus of excitation and can yield *attractor-like* category-specific spatiotemporal patterns (e.g., resonant responses) determined by the OWM dynamics rather than by the mapping alone.

We monitor changes in the OWM system energy against an adaptive saliency threshold T_{adapt} . The OWM system energy is defined as the sum of the square of each neuron value at each time step Δt . The T_{adapt} is computed using an energy-based approach:

$$T_{\text{adapt}} = \mu + 2\sigma(1 + \alpha \cdot \text{trend}),$$

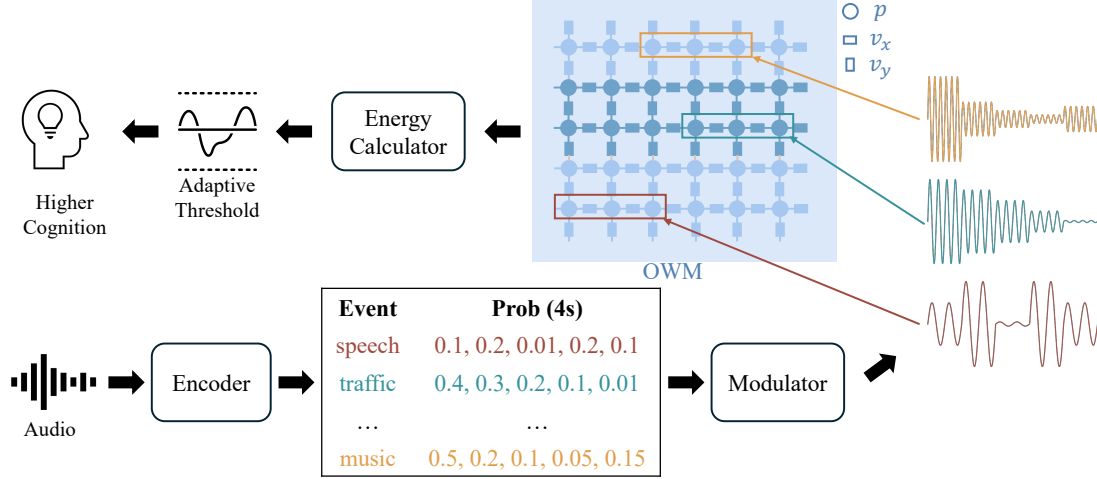


Figure 2. **Overview of the NAACA (NeuroAuditory Attentive Cognitive Architecture).** Audio is segmented into sliding windows and mapped by a pretrained encoder to auditory object probability trajectories, which drive frequency-specific oscillatory inputs on OWM grids. OWM is a 2D neural network with primary (p) and velocity (v) neurons, parameterized by wave propagation speed c and damping k , where c follows a stripe-shaped binary pattern (dark/light blue). A gate-opening decision is made from OWM energy fluctuations relative to an adaptive threshold, and salient audio segments are routed to a higher-level cognition module for semantic interpretation.

where μ and σ are the running mean and standard deviation of energy-derived drift metrics over a sliding window of $W = 20$ samples, and the trend factor adjusts for temporal patterns in the data. The final detection decision employs persistence filtering to ensure robust saliency detection while minimizing false positives, after which the detected segments are forwarded to a higher-level cognitive module for semantic interpretation. The whole procession is shown in Algorithm C.2.

2.2. OWM Formulation

OWM is a 2D recurrent field model defined on a $G \times G$ lattice whose dynamics are analyzed and designed based on a discretized damped wave system. Its hidden state $\mathbf{h}_t \in \mathbb{R}^{G^2}$ stores a primary field $p(x, y, t)$ whose activity propagates via a structured spatial operator $\mathcal{A}(\cdot)$ with wave-like coupling. The update combines (i) temporal recurrence through \mathbf{h}_{t-1} for memory and (ii) spatial recurrence through $\mathcal{A}(\cdot)$ for lateral propagation. An intermediate velocity-like variable $\mathbf{v}(x, y, t) = (v_x, v_y)$ mediates directional flow. Parameters c , k^p , and k^v control propagation speed and damping. This bio-inspired design is analogous to membrane-potential storage (p) and axonal/dendritic transport (\mathbf{v}) in cortical sheet.

2.2.1. CONSTRUCTING THE STRUCTURED SPATIAL OPERATOR $\mathcal{A}(\cdot)$ OF OWM

Wave System Foundations. We consider the OWM as a two-dimensional oscillatory system governed by a damped wave equation in the first-order velocity–pressure formula-

tion:

$$\begin{aligned} \frac{\partial p}{\partial t} + k^p(x, y) p &= -c^2(x, y) \nabla \cdot \mathbf{v} + S(x, y, t), \\ \frac{\partial \mathbf{v}}{\partial t} + k^v(x, y) \mathbf{v} &= -\nabla p, \end{aligned} \quad (2)$$

where $c(x, y)$ is the spatially varying wave propagation speed (time-independent), $k^p(x, y)$ and $k^v(x, y)$ are pressure and velocity damping coefficients (time-independent), and $S(x, y, t)$ is given by Eq. 1, and $\nabla = [\partial/\partial x, \partial/\partial y]^T$ is the gradient operator, and $\nabla \cdot \mathbf{v} = \partial v_x/\partial x + \partial v_y/\partial y$ is the divergence operator.

Discretizing Eqs. 2 over a two-dimensional lattice yields:

$$\begin{aligned} p(x, y, t + \Delta t) &= [1 - \Delta t k^p(x, y)] p(x, y, t) \\ &\quad - \Delta t c^2(x, y) \nabla \cdot \mathbf{v}(x, y, t) + \Delta t S(x, y, t), \\ \mathbf{v}(x, y, t + \Delta t) &= [1 - \Delta t k^v(x, y)] \mathbf{v}(x, y, t) - \Delta t \nabla p(x, y, t). \end{aligned} \quad (3)$$

where Δt is the time step.

This discrete update system exhibits frequency-selective response through spatial phase differentiation, as characterized by Theorem 2.1.

Theorem 2.1 (Frequency-Selective Phase Response). *Under multi-frequency excitation $S(x, y, t) = \sum_i S_i(x, y, t)$ with components at frequencies $\{\omega_i\}$, the steady-state pressure field exhibits frequency-dependent phase delay:*

$$p(x, y, t) \approx \sum_i A_i(x, y) S_i(x, y, t - \tau_i(x, y)), \quad (4)$$

where the phase delay is determined by input-resonance

mismatch:

$$\tau_i(x, y) \propto \arctan \left(\frac{\omega_i - \omega_{\text{res}}(x, y)}{k^p(x, y) + k^v(x, y)} \right), \quad (5)$$

with $\omega_{\text{res}}(x, y)$ the local resonance frequency, $A_i(x, y)$ the amplitude response, and $\tau_i(x, y)$ the temporal phase delay at location (x, y) for input frequency ω_i . This establishes frequency selectivity via phase: different frequencies produce distinct phase delays, enabling spatial encoding of spectral content.

Proof. See Appendix D. \square

System Energy Measurement. Different spatial locations within the OWM lattice are associated with distinct attractors corresponding to specific sound event categories, enabling the system to localize and track the perceptual saliency of diverse auditory sources. Since each event is characterized by its own modulatory input frequency, the local wave propagation speed $c(x, y)$ must also vary across space. This spatial dependence ensures that the eigenfrequency structure reflects the diversity of sound-driven dynamics in the system.

However, while local eigenmodes are essential for modeling event-specific resonances, pattern drift detection and memory-related computations require more than analyzing these modes in isolation. Instead, the collective behavior of the system must be captured in terms of a global state variable. To this end, we define the total system energy, which aggregates pressure and velocity contributions across the lattice. This energy-based representation not only reflects the ongoing dynamics of the OWM but also forms the key computational signal for detection and optimality analysis.

The total energy of the OWM system in discrete form is

$$E(t) = \frac{1}{2} \sum_{i,j} [p_{i,j}^2(t) + v_{x,i,j}^2(t) + v_{y,i,j}^2(t)], \quad (6)$$

where the terms correspond to kinetic energy due to coupling, potential energy due to stiffness. For the purpose of analyzing and designing the internal structure and parameters, we approximate the 2D lattice as a continuous medium. The energy then becomes

$$E(t) = \iint \frac{1}{2} [p^2(x, y, t) + v_x^2(x, y, t) + v_y^2(x, y, t)] dx dy, \quad (7)$$

which will be used in the following calculations and theorem proofs as the basis for sensitivity and optimality analyses.

2.2.2. TOPOLOGICAL ORGANIZATION WITH HIGH SENSITIVITY TO SALIENCE

To analyze the role of topological organization and its sensitivity to drift, we first reformulate the governing first-order velocity–pressure system in a more compact representation

in Theorem 2.2. This reformulation exposes the effective damping and restoring mechanisms and serves as a foundation for later connecting topological behavior with energy dynamics.

Theorem 2.2 (Equivalence of First-Order System to Second-Order Damped Wave Equation). *The first-order pressure-velocity system Eqs. 2 is equivalent to the second-order damped wave equation*

$$\frac{\partial^2 p}{\partial t^2} + \gamma \frac{\partial p}{\partial t} + k^v k^p \cdot p = \nabla \cdot (c^2 \nabla p) + (k^v \mathbf{v} + \nabla p) \cdot \nabla c^2 + \left(k^v + \frac{\partial}{\partial t} \right) S, \quad (8)$$

with effective damping coefficient $\gamma = k^p + k^v$, restoring force coefficient $\mu = k^v k^p$, and modified source term $S_{\text{eff}} = (k^v + \partial/\partial t)S$. The term $(k^v \mathbf{v} + \nabla p) \cdot \nabla c^2$ represents spatial coupling from wave speed variation.

Proof. See Appendix E. \square

Building on this structural equivalence, we next derive the explicit energy evolution law, which highlights how the wave propagation speed c governs energy redistribution and thereby influences stability and drift sensitivity.

Theorem 2.3 (Energy Evolution). *When $p(x, y, t)$ and v_x, v_y are governed by Eqs. 2 under periodic boundaries. The energy evolution is given by*

$$\frac{dE}{dt} = \iint [pS - k^p p^2 - k^v |\mathbf{v}|^2 - (c^2 - 1)p \nabla \cdot \mathbf{v}] dx dy. \quad (9)$$

Proof. See Appendix F. \square

The injection term $\iint pS dx dy$ in Eq. (9) governs detection sensitivity. To maximize sensitivity, $c(x, y)$ must create frequency-selective spatial differentiation where different auditory events drive spatially distinct activity patterns through varying local resonance frequencies (Lakatos et al., 2016), while enabling coherent coupling between slow maintenance dynamics and fast encoding transients (Lundqvist et al., 2018). A striped structure $c(x, y) = c(y)$ with Bragg-matched periodicity (Kushwaha et al., 1993) achieves both objectives: coherent reflections create slow-propagating modes supporting maintenance oscillations, which interact with encoding transients via phase-dependent energy injection (Lundqvist et al., 2018).

Theorem 2.4 (Striped Pattern Optimality). *Under amplitude constraint $|\delta c^2| \leq A$, the Bragg-resonant striped square wave $\delta c^2(x, y) = A \cdot \text{sgn}[\cos(2\pi q_0 y/L_y)]$ maximizes modal coupling strength $|C_{(m,n),(m,n+q_0)}|$ and spatial frequency differentiation for salience detection at mode separation q_0 .*

Proof. See Appendix G.4. \square

The Bragg-matched structure supports multi-timescale dynamics through slow-propagating coherent modes, enabling interactions between maintenance and encoding processes analogous to cross-frequency coupling in neural working memory (Lundqvist et al., 2018).

3. Experiment

Dataset. We evaluate NAACA using XD-Violence (Wu et al., 2020) and Urban Soundscapes of the World (USoW) (De Coensel et al., 2023; 2017). For the multi-modal XD-Violence dataset, we utilize only the audio track to benchmark our auditory drift detection. As NAACA is training-free, we evaluate exclusively on the 500-sample test set, leveraging its high-granularity, frame-level labels to assess temporal boundary precision. To complement the stylized movie audio of XD-Violence, we use USoW to simulate real-world surveillance. USoW provides unlabeled, high-quality 4-channel ambisonics of stable urban environments punctuated by rare events. This serves as our primary benchmark for qualitative analysis, demonstrating OWM’s robustness against ambient noise and its interpretability in detecting meaningful sonic shifts.

Measurements. We evaluate our framework using a mix of ground-truth metrics and efficiency proxies. For XD-Violence, we utilize frame-level annotations to report Average Precision (AP), assessing the global ability to distinguish violent anomalies from background noise. To measure the practical utility of our attention-gating across both datasets, we define Computational Time Saved. This metric quantifies the reduction in inference overhead achieved by invoking the ALM only at detected drift points versus exhaustive, continuous processing. By reporting the fraction of detected drifts and visualizing the OWM’s internal energy states, we demonstrate the architecture’s ability to capture salient shifts while minimizing the computational burden of long-form audio understanding.

Baselines. We evaluate our framework against three categories: (i) *Exhaustive ALM Inference*, an internal ablation where AudioQwen processes the full stream without gating to isolate the efficiency gains of our mechanism; (ii) *Supervised audio-only models*, including HL-Net (Wu et al., 2020) and AVadCLIP (Wu et al., 2025); and (iii) *Video-only frameworks*, covering supervised models such as S3R (Wu et al., 2022), VadCLIP (Wu et al., 2024) and Holmes-VAU (Zhang et al., 2025), as well as zero-shot approaches like TRACE (Siddiqui et al., 2025). Crucially, unlike our training-free approach, these baseline models require domain-specific fine-tuning, extra training on domain knowledge, or auxiliary attention fusion layers.

Implementation Details. Audio is processed in 4 s windows using PANN (Kong et al., 2020) features (527 classes). OWM uses a 64×64 grid with $dt = 0.01$, $k_p = k_v = 10$, and striped $c(y) \in [0.1, 70]$ from Theorem 2.1. AudioQwen (Chu et al., 2023) provides semantic interpretation. Full details in Appendix B.

4. Results and Discussion

4.1. Quantitative Evaluation on XD-Violence

Performance Success. The primary goal of our framework is to achieve competitive audio-only performance using general-purpose models without the need for expensive training from scratch or domain-specific fine-tuning. Task-specific models often incur prohibitive costs in data collection and labeling, and video-based systems are frequently limited by environmental factors such as poor lighting or occlusions; in contrast, audio-based models offer a more resilient alternative that is less sensitive to such conditions.

Table 1. Average Precision (AP) on XD-Violence, grouped by baseline type.

Method	Training	Zero-shot	Modality	AP (%)
Audio Qwen	✓		Audio	53.50
HL-Net [†]	✓		Audio	60.50
AVadCLIP [†]	✓		Audio	52.51
S3R*	✓		Video	80.26
VadCLIP*	✓		Video	84.51
Holmes-VAU*	✓		Video	87.68
Zero-shot CLIP*		✓	Video	17.83
TRACE [†]	✓	✓	Video	83.67
NAACA		✓	Audio	70.60

*Results reported in Holmes-VAU paper.

[†]Results reported in their respective papers.

As shown in Table 1, we utilize Audio Qwen as our operational baseline for classifying the audio stream. While video-only models generally outperform audio-only models on the XD-Violence task due to the explicit visual cues available in action-movie-style data, our results demonstrate that audio remains a highly informative modality when processed correctly.

Our NAACA achieves an AP of 70.60%. This represents a substantial 17.1% absolute improvement over the baseline Audio Qwen (53.50% AP). Notably, our training-free approach significantly outperforms supervised audio-only models like AVadCLIP (52.51% AP) and HL-Net (60.50% AP). While video-based methods like TRACE show strong results (83.67% AP), it is important to note that TRACE is not a pure zero-shot solution in the same sense as our framework; despite being marked as zero-shot in Table 1, it relies heavily on the training of an additional temporal encoder and a cross-attention fusion layer, which we consider a form of domain adaptation. In contrast, our framework provides a specialized, training-free solution for scenarios where visual data is unavailable or where computational resources preclude the training of auxiliary fusion modules.

The Audio-Video Gap (Failure Case Analysis). Despite NAACA’s competitive performance, a gap remains between audio-only frameworks and supervised video models. As illustrated in the confusion matrix in Fig. 3, this performance ceiling is largely attributable to inherent acoustic ambiguity and dataset constraints. The *Abuse* category suffers the highest misclassification, with 45.0% of cases identified as *Fighting*. This is primarily due to the extreme scarcity of data—comprising only 11 samples—and the high semantic overlap of acoustic cues like shouting. Similarly, *Shooting* exhibits a 19.0% confusion rate with *Fighting*. These events frequently co-occur in real-world scenarios, and without visual cues to identify firearms, the model defaults to the broader acoustic signature of a physical altercation. Furthermore, categories such as *Explosions* and *Car Accidents* rely heavily on visual markers like fire or debris for definitive classification. While 14.7% of *Fighting* samples are missed entirely (appearing in the None/Miss category), only 0.7% are confused with *Explosions*, though sudden, high-energy acoustic transients from both categories can be acoustically similar without visual context.





Figure 3. Confusion matrix on the XD-Violence test set audio track. Significant overlaps between *Abuse*, *Shooting*, and *Fighting* reflect acoustic ambiguities and event co-occurrence. The misclassification of *Fighting* as *Explosions* highlights the reliance on visual cues for high-energy transient events.

4.2. Qualitative Case Study on USoW


To qualitatively assess the behavior of our framework, we highlight three illustrative cases of auditory pattern drift. Specifically, we consider: (1) the detection of salient novel events that shifts the auditory context, (2) robustness to transient silences or pauses in ongoing sound streams, and (3) sensitivity to subcategory-level substitutions within a broader sound category (e.g., different types of musical instruments within the “music” class). These cases provide concrete insights into how the OWM detects and distinguishes different forms of drift beyond low-level acoustic fluctuations.

4.2.1. DETECTION OF SALIENT NOVEL EVENTS

We illustrate this case with two representative recordings in which the most salient novel event occurs near the end of the segment. In both recordings, the background soundscape is relatively stable: the Montreal recording features continuous traffic flow and bird chirping, while the Berlin recording contains crowd conversations, birds, and fountain noise. This provides a clear contrast when a new source emerges late in the sequence, allowing us to assess how well each method captures such onsets.



(a) R0002 (Place d’Armes, Montreal): Car engine onset at 53 s; OWM *p*-field activation near (30, 52).



(b) R0056 (Alexanderplatz, Berlin): Bagpipe onset at 49 s; OWM *p*-field activation near (25, 22).


Figure 4. OWM detection of novel events. Mel-spectrograms with car engine (a) and bagpipe (b) onsets. OWM outputs (cyan = energy change, red = adaptive threshold, *p*-field states shown). Vertical dashed lines mark detected drifts.

As shown in Fig. 4, OWM successfully detects late-arriving novel events with high precision. While traditional embedding-based approaches often suffer from spurious detections due to spectral variance, OWM localizes the onset stably (e.g., in Example R0002, Fig. 4a). This stability arises because the OWM *p*-field internal states concentrate energy precisely around the true change point rather than scattering across the segment. This behavior demonstrates OWM’s robustness in distinguishing genuine novel events from background variability.


4.2.2. ROBUSTNESS TO TRANSIENT PAUSES

We next consider recordings where the salient sound events include natural pauses or interruptions. Such cases are challenging because detectors may mistake short gaps within an ongoing event for the onset of new events. The first recording (Fig. 5a) was captured in a dense traffic environ-

ment, while the second (Fig. 5b) comes from a lively public square during a music festival. In both recordings, prominent sources exhibit intermittent activity, providing a useful testbed for evaluating robustness to transient pauses.



(a) R0037 (Johnston Road, Hong Kong): A baby cry with short pauses. The baseline repeatedly flags pauses, while OWM registers one event at 33 s.



(b) R0016 (Quincy Market, Boston): Festival scene with speech followed by applause. The baseline reports near-continuous changes, while OWM yields a single detection.

Figure 5. OWM is robust to transient pauses. Mel-spectrogram of OWM output. Vertical dashed lines mark detected drifts: cyan indicates energy change, and red indicates the adaptive threshold.

As shown in Fig. 5, OWM successfully identifies the main events while maintaining stability through spectral fluctuations. While standard detection methods often produce frequent spurious triggers whenever short silences occur (as seen in the middle and end portions of Example R0016), OWM effectively consolidates these interruptions. This demonstrates its ability to maintain a stable event representation and avoid over-segmentation despite transient pauses in the acoustic signal.

4.2.3. SUBCATEGORY TRANSITION SENSITIVITY

We further analyze cases where acoustic variation arises within a subcategory of ongoing sounds rather than from the introduction of a completely new source. In this example (Fig. 6), recorded in Square Phillips (Montreal), the background consists of street traffic mixed with music containing two main instruments: a hi-hat at higher frequencies and a slower kick drum. The interplay of these instruments produces several subcategory-level shifts.

Specifically, the hi-hat drops out at 21 s, leaving only the kick drum; it reappears at 32 s and pauses again at 38 s. OWM successfully detects each of these subtle transitions, capturing the precise moments of instrument entry and exit.

At 42 s, a distinct car horn emerges, which OWM identifies as a salient novel source. These results demonstrate OWM’s ability to capture fine-grained subcategory-level drift while simultaneously maintaining high sensitivity to distinct novel events.




Figure 6. OWM sensitivity to fine-grained salient transitions (Example R0010, Square Phillips, Montreal). The segment features alternating hi-hat and kick drum patterns, with a distinct car horn at 42 s that OWM detects as a novel event alongside the subcategory drift. Vertical dashed red lines mark detected drift points. The cyan curve represents energy change, the red curve shows the adaptive threshold, and detected events are highlighted at the change points.

4.3. Spectral Analysis of p -field Dynamics via FFT

We computed Fast Fourier Transforms (FFTs) of p -field activity, sampling internal states every second to match the sliding window stride (Subsection 3). With $dt = 0.01$, one second equaled 100 time steps. For each p neuron, we extracted the dominant frequency (maximal FFT amplitude) to construct frequency maps (Fig. 7, Examples R0016 and R0056). To suppress numerical noise, only neurons above the 75th percentile of variance were retained. Additional results appear in Appendix J.

The analysis revealed spatially clustered oscillatory activity rather than uniform grid activation, with dominant frequencies limited to 0–50 Hz by the Nyquist bound ($dt = 0.01$ s). Distinct frequency bands aligned with canonical neural regimes: during stable background periods, both examples exhibited sustained β -band activity (15–30 Hz), consistent with its theorized role in working memory maintenance (Lundqvist et al., 2018). After drift onset, Example R0016 shifted toward γ -band activity (30–50 Hz), reflecting rapid encoding of salient auditory input (applause, cheering). In Example R0056, the post-drift segment with emerging bagpipe sounds showed subsets of neurons entering the γ range, while pronounced β -band oscillations persisted, suggesting mixed maintenance and encoding dynamics. Spatially, both examples exhibited a redistribution of activity: channels along the upper grid boundary ($Y=1$ row), previously speech-related, showed reduced activity post-drift, whereas deeper clusters ($Y \approx 20$) became strongly engaged, consistent with recruiting new resources for encoding applause, cheering, and musical instruments. These patterns underscore that OWM reallocates oscillatory dynamics to novel salient sources rather than sustaining prior speech inputs.

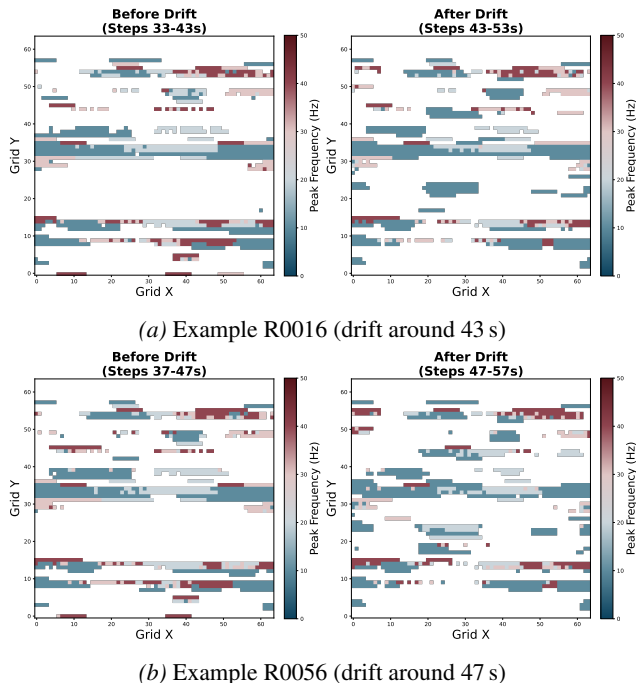


Figure 7. **Temporal frequency analysis around drift detection events.** Frequency distributions in active p neurons during 10 s before (left) and after (right) drift onset. Only neurons above the 75th percentile activity threshold are shown.

Supplementary animations show the real-time evolution of p -field activity for each example.

In contrast, α -band activity (8–12 Hz) remained weak with no systematic changes across drift events, precluding confirmation of links between elevated α power and attentional lapses (Lakatos et al., 2016; Kasten et al., 2024). Similarly, θ -band activity (4–8 Hz) was sparse and failed to form robust clusters, despite prior reports of θ entrainment supporting auditory working memory (Albouy et al., 2017; Bonetti et al., 2024).

In sum, the OWM reallocates frequency-specific oscillatory clusters, similar to cortical γ -band encoding and β -band maintenance, to represent salient auditory transitions within its bio-physical constraints.

4.4. Cognitive Resource Efficiency

Deploying ALMs for long-form audio requires balancing accuracy with computational cost. While exhaustive sliding-window processing addresses attention dilution, it requires 57 ALM invocations per 60 s clip, which is prohibitive for real-time use. NAACA optimizes this via OWM-based gating, forwarding only salient segments. We measure efficiency using the Time Sent Ratio.

As shown in Fig. 8, median ratios for XD-Violence and USoW are 0.597 and 0.650, respectively. This signifies




Figure 8. **Time sent ratios for XD-Violence and USoW datasets.** Violin plots with box plots and scatter points show the fraction of audio forwarded to the ALM after OWM drift detection. Both datasets exhibit similar distributions (medians: 0.597 and 0.650), demonstrating that NAACA consistently processes only 60% of audio duration, substantially reducing computational cost while preserving detection accuracy.

a 40% reduction in ALM invocations (from ~ 57 to ~ 34). Notably, this cost reduction is paired with a 17.1% improvement in AP (70.60% vs. 53.50% baseline). By focusing on drift-enriched windows, OWM provides the ALM with contexts where transitions are concentrated, shifting the Pareto frontier of accuracy and efficiency.

Scatter points in Fig. 8 demonstrate that NAACA’s cost scales adaptively: ratios near 1.0 indicate high acoustic complexity, while stable soundscapes yield ratios of 0.3–0.4. This adaptive allocation spends the computational budget proportionally to information density. Despite different acoustic characteristics, OWM maintains stable gating across both datasets, positioning NAACA as a practical solution for resource-constrained monitoring.

5. Conclusion

We introduced NAACA, a neuro-inspired framework for ALM understanding enhancement, with OWM as its core working-memory component. Our approach combines a wave-based recurrent field model with an energy-driven drift detection mechanism that adaptively reallocates attention without long-term historical data or offline pretraining. We proved that binary and striped wave-speed distributions optimize drift sensitivity, and demonstrated through urban soundscape experiments that OWM reliably captures salient novel events, resists transient pauses, and subcategory saliency shifts more effectively than similarity-based baselines. OWM’s oscillatory dynamics align with cortical working memory, underscoring its biological plausibility and interpretability. These results establish OWM as a computationally efficient, neuro-inspired foundation for extending long-context reasoning in multimodal systems.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Aggarwal, C. C., Hinneburg, A., and Keim, D. A. On the surprising behavior of distance metrics in high dimensional space. In *International conference on database theory*, pp. 420–434. Springer, 2001.
- Albouy, P., Weiss, A., Baillet, S., and Zatorre, R. J. Selective entrainment of theta oscillations in the dorsal stream causally enhances auditory working memory performance. *Neuron*, 94(1):193–206, 2017.
- Bai, J., Yin, H., Wang, M., Shi, D., Gan, W.-S., Chen, J., and Rahardja, S. Audiolog: Llms-powered long audio logging with hybrid token-semantic contrastive learning. In *2024 IEEE International Conference on Multimedia and Expo (ICME)*, pp. 1–6. IEEE, 2024.
- Beyer, K., Goldstein, J., Ramakrishnan, R., and Shaft, U. When is “nearest neighbor” meaningful? In *International conference on database theory*, pp. 217–235. Springer, 1999.
- Bonetti, L., Fernández-Rubio, G., Carlomagno, F., Dietz, M., Pantazis, D., Vuust, P., and Kringelbach, M. Spatiotemporal brain hierarchies of auditory memory recognition and predictive coding. *Nature Communications*, 15(1):4313, 2024.
- Brennan, C. and Proekt, A. Attractor dynamics with activity-dependent plasticity capture human working memory across time scales. *Communications psychology*, 1(1):28, 2023.
- Chan, T.-H. H., Jiang, S. H.-C., Wu, T., and Zhao, M. Online clustering with nearly optimal consistency. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=NA2vUMaMOM>.
- Chu, Y., Xu, J., Zhou, X., Yang, Q., Zhang, S., Yan, Z., Zhou, C., and Zhou, J. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*, 2023.
- Corbetta, M. and Shulman, G. L. Control of goal-directed and stimulus-driven attention in the brain. *Nature reviews neuroscience*, 3(3):201–215, 2002.
- De Coensel, B., Sun, K., and Botteldooren, D. Urban soundscapes of the world: Selection and reproduction of urban acoustic environments with soundscape in mind. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, volume 255, pp. 5407–5413. Institute of Noise Control Engineering, 2017.
- De Coensel, B., Botteldooren, D., Kang, S., and Van Renterghem, T. *Urban soundscapes of the world*. zenodo: <https://zenodo.org/records/10106181>, 2023. doi: 10.5281/zenodo.10106180.
- Friston, K. The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7):293–301, 2009.
- Greco, S., Vacchetti, B., Apiletti, D., and Cerquitelli, T. Unsupervised concept drift detection from deep learning representations in real-time. *IEEE Transactions on Knowledge and Data Engineering*, 2025.
- He, B., Li, H., Jang, Y. K., Jia, M., Cao, X., Shah, A., Shrivastava, A., and Lim, S.-N. Ma-Imm: Memory-augmented large multimodal model for long-term video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13504–13514, 2024.
- Kasten, F. H., Busson, Q., and Zoefel, B. Opposing neural processing modes alternate rhythmically during sustained auditory attention. *Communications Biology*, 7(1):1125, 2024.
- Kayser, C., Petkov, C. I., Lippert, M., and Logothetis, N. K. Mechanisms for allocating auditory attention: an auditory saliency map. *Current biology*, 15(21):1943–1947, 2005.
- Kong, Q., Cao, Y., Iqbal, T., Wang, Y., Wang, W., and Plumbley, M. D. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2880–2894, 2020. doi: 10.1109/TASLP.2020.3030497.
- Kushwaha, M. S., Halevi, P., Dobrzynski, L., and Djafari-Rouhani, B. Acoustic band structure of periodic elastic composites. *Physical review letters*, 71(13):2022, 1993.
- Lakatos, P., Barczak, A., Neymotin, S. A., McGinnis, T., Ross, D., Javitt, D. C., and O’Connell, M. N. Global dynamics of selective attention and its lapses in primary auditory cortex. *Nature neuroscience*, 19(12):1707–1717, 2016.
- Lundqvist, M., Herman, P., Warden, M. R., Brincat, S. L., and Miller, E. K. Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nature communications*, 9(1):394, 2018.

- 495 Qian, R., Ding, S., Dong, X., Zhang, P., Zang, Y., Cao, Y.,
 496 Lin, D., and Wang, J. Dispider: Enabling video llms with
 497 active real-time interaction via disentangled perception,
 498 decision, and reaction. In *Proceedings of the Computer
 499 Vision and Pattern Recognition Conference*, pp. 24045–
 500 24055, 2025.
- 501 Rabanser, S., Günnemann, S., and Lipton, Z. Failing loudly:
 502 An empirical study of methods for detecting dataset shift.
 503 *Advances in Neural Information Processing Systems*, 32,
 504 2019.
- 505 Siddiqui, Y. A., Usmani, S., Tariq, U., Shamsi, J. A., and
 506 Khan, M. B. Traces: Temporal recall with contextual
 507 embeddings for real-time video anomaly detection. *arXiv
 508 preprint arXiv:2511.00580*, 2025.
- 509 Wan, K., Liang, Y., and Yoon, S. Online drift detection with
 510 maximum concept discrepancy. In *Proceedings of the
 511 30th ACM SIGKDD Conference on Knowledge Discovery
 512 and Data Mining*, pp. 2924–2935, 2024.
- 513 Winkler, I., Denham, S. L., and Nelken, I. Modeling the
 514 auditory scene: predictive regularity representations and
 515 perceptual objects. *Trends in cognitive sciences*, 13(12):
 516 532–540, 2009.
- 517 Wu, J.-C., Hsieh, H.-Y., Chen, D.-J., Fuh, C.-S., and Liu,
 518 T.-L. Self-supervised sparse representation for video
 519 anomaly detection. In *European Conference on Computer
 520 Vision*, pp. 729–745. Springer, 2022.
- 521 Wu, P., Liu, J., Shi, Y., Sun, Y., Shao, F., Wu, Z., and Yang,
 522 Z. Not only look, but also listen: Learning multimodal
 523 violence detection under weak supervision. In *European
 524 conference on computer vision*, pp. 322–339. Springer,
 525 2020.
- 526 Wu, P., Zhou, X., Pang, G., Zhou, L., Yan, Q., Wang, P., and
 527 Zhang, Y. Vadclip: Adapting vision-language models for
 528 weakly supervised video anomaly detection. In *Proceed-
 529 ings of the AAAI Conference on Artificial Intelligence*,
 530 volume 38, pp. 6074–6082, 2024.
- 531 Wu, P., Su, W., Pang, G., Sun, Y., Yan, Q., Wang, P.,
 532 and Zhang, Y. Avadclip: Audio-visual collaboration
 533 for robust video anomaly detection. *arXiv preprint
 534 arXiv:2504.04495*, 2025.
- 535 Wu, X., Lu, H., Li, K., Wu, Z., Liu, X., and Meng, H.
 536 Hiformer: Sequence modeling networks with hierarchical
 537 attention mechanisms. *IEEE/ACM Transactions on Audio,
 538 Speech, and Language Processing*, 31:3993–4003, 2023.
- 539 Zhang, H., Xu, X., Wang, X., Zuo, J., Huang, X., Gao, C.,
 540 Zhang, S., Yu, L., and Sang, N. Holmes-vau: Towards
 541 long-term video anomaly understanding at any granular-
 542 ity. In *Proceedings of the Computer Vision and Pattern
 543 Recognition Conference*, pp. 13843–13853, 2025.
- 544
 545
 546
 547
 548
 549

Appendix

A. Related Work

Auditory selective attention and temporal pattern drift have been studied both in neuroscience, where oscillatory dynamics are linked to memory and attentional control, and in machine learning, where attention bottlenecks in audio language models (ALMs) constrain long-sequence processing. Insights from neuroscience on tracking behaviorally relevant inputs and adapting to changes over time motivate analogous strategies for ALMs operating on long audio streams. In our setting, we use *temporal drift* to denote a shift in the audio stream’s statistics or in its task-relevant representations (e.g., embedding or class-probability trajectories) across time. This is closely related to *concept drift* in machine learning, which refers to changes in the underlying data-generating process that alter the mapping between inputs and the semantics or labels of interest. Thus, detecting temporal drift in representations provides an online, label-free proxy for potential concept drift and can trigger reallocation of limited inference resources.

A.1. Oscillatory Dynamics in Memory and Attention

One prevailing hypothesis is that working memory is supported by discrete neural activity configurations, often described as attractor states (Brennan & Proekt, 2023). However, attractor dynamics alone may fail to explain working memory across multiple timescales, motivating models that couple attractor-like dynamics with activity-dependent plasticity (Brennan & Proekt, 2023). Within this dynamical view, neural oscillations provide a complementary mechanism for coordinating information maintenance and prioritization over time. Neural oscillations in specific frequency bands have been shown to play a central role in memory, selective attention, and sensitivity to temporal drift. For example, (Lundqvist et al., 2018) demonstrated that working memory tasks involve non-stationary dynamics, with gamma bursts during encoding and beta bursts during maintenance. Similarly, selective entrainment of theta oscillations has been shown to enhance auditory working memory performance (Albouy et al., 2017; Bonetti et al., 2024). In contrast, high alpha-band activity has been associated with increased error rates and reduced auditory attention (Lakatos et al., 2016; Kasten et al., 2024).

A.2. Attention Limitations and Concept Drift Detection

Inspired by these oscillatory mechanisms, one can view attention in ALMs as a resource that must be selectively and dynamically allocated in response to temporal pattern drift rather than distributed uniformly across the input. Although recent advances in ALMs have enabled significant progress in audio understanding, long-form reasoning remains limited by restricted attention span, motivating efforts to extend context length (Wu et al., 2023; He et al., 2024; Bai et al., 2024). Most existing solutions require retraining or fine-tuning, which is computationally costly and inflexible. A complementary perspective is to frame attention allocation as a drift detection problem, where established methods from the machine learning literature may provide efficient and adaptive alternatives.

In real-world audio monitoring, *salient* events are inherently rare and context dependent, and many of the situations that make them important are only evident at run time. As a result, it is generally impractical to construct exhaustive supervised datasets that cover all relevant foreground events and background conditions, and to train a dedicated detector for each deployment scenario. Instead, the scope can be narrowed by operating in a task-relevant *semantic* space: a high-dimensional auditory representation learned from large corpora that captures human-recognizable sound structure and supports downstream reasoning. Within this representation space, salience can be operationalized as online distributional change, motivating unsupervised drift detection methods that trigger selective processing without requiring explicit event labels.

High-dimensional semantic embedding challenges. Accordingly, our setting calls for online drift detection in semantic audio representations, such as class-probability trajectories or embedding sequences produced by pretrained audio recognition models. By operating in this high-dimensional representation space learned from large corpora, changes in acoustic context manifest as distributional shifts over time, providing a practical signal for unsupervised drift detection without requiring explicit event labels. Naive statistical approaches, such as tracking Euclidean distances from exponential moving averages, face the curse of dimensionality: as dimensions increase, distance metrics lose discriminative power as all points become approximately equidistant (Aggarwal et al., 2001; Beyer et al., 1999). A key difficulty in high-dimensional drift detection is that the same change magnitude can correspond to qualitatively different shifts. For example, consider two consecutive class-probability vectors with identical L_2 distance: in one case, probability mass moves coherently between a small set of semantically related classes (e.g., *gunshot* and *explosion*); in another case, the same L_2 distance arises from small, diffuse fluctuations spread across many unrelated classes. A single threshold on $\|p_{t+1} - p_t\|_2$ cannot reliably distinguish these

cases, motivating drift measures that exploit structure in the semantic space rather than magnitude alone.

Historical data and training requirements. Cluster-based approaches (Chan et al., 2025) assume sufficient samples per category to maintain stable cluster statistics, which is unrealistic for heterogeneous urban soundscapes where salient events (e.g., sirens, screams) are rare and unpredictable. MCD-DD (Wan et al., 2024) employs contrastive learning on encoder representations to compute maximum concept discrepancy, but requires maintaining reference distributions from extensive historical windows (typically >100 samples) and incurs significant computational overhead from pairwise similarity computation. DriftLens (Greco et al., 2025) detects drift from deep representations in real time, but assumes an *offline* instantiation phase that estimates a fixed reference (baseline) embedding distribution and detection thresholds from historical/training-era data. A different setting is training-free stream monitoring, where semantic representations are extracted from pretrained audio encoders but access to encoder training data or a dedicated in-domain reference corpus may be unavailable, making offline baseline construction less straightforward.

Deployment constraints. Real-time audio monitoring, such as public safety surveillance or environmental monitoring, imposes strict constraints: (1) *No offline training*: Systems must operate on new deployments without domain-specific fine-tuning; (2) *No big historical buffering*: Embedded devices cannot store hours of high-dimensional embeddings; (3) *Cold-start capability*: Detection must work immediately without warm-up periods; (4) *Non-stationary environments*: Acoustic distributions evolve (morning traffic → evening crowds), rendering historical baselines obsolete.

Table A.1 summarizes these requirements. Our OWM-based approach is designed to satisfy them by maintaining a compact online state via oscillatory dynamics, which summarizes recent context without storing long histories of high-dimensional embeddings. Here training-free refers to requiring no additional fine-tuning or domain adaptation when deployed in a new environment (given a fixed pretrained audio encoder/ALM), and no long historical data accumulation refers to not requiring an offline reference corpus or baseline calibration phase prior to monitoring.

Table A.1. Requirements comparison of drift detection approaches for audio-semantic streams. *Training-free*: No offline training or fine-tuning required. *Online*: Processes streaming data without batch accumulation. *No long historical data*: Does not require long-term buffering (>50 samples) of past embeddings or statistics. *High-dim*: Effective on >500 dimensional semantic embeddings without distance concentration issues. *Topology-preserving*: Maintains structured spatial relationships in semantic representation space, not just raw semantic features. The symbol ~ indicates partial support or conditional availability depending on configuration parameters.

Method	Training-free	Online detection	No long historical data	High-dim (>500D)	Topology -preserving
Statistical (EMA/Variance)	✓	✓	✓	×	×
MCD-DD (Wan et al., 2024)	×	✓	×	✓	×
DriftLens (Greco et al., 2025)	×	~	×	✓	×
Cluster-based (Chan et al., 2025)	×	✓	×	~	×
OWM (ours)	✓	✓	✓	✓	✓

A.3. Selective Attention in Video-Language Models

The challenge of redundant long-context processing is a central theme in recent Video Language Models (VLMs). Dispider (Qian et al., 2025) introduces a disentangled perception framework where a “Decision” module monitors coarse-grained video to proactively trigger a “Reaction” module only during relevant interaction moments. Similarly, Holmes-VAU (Zhang et al., 2025) utilizes an anomaly-focused Temporal Sampler (ATS) that leverages a lightweight detector to generate anomaly scores, allowing the VLM to selectively aggregate salient segments while filtering uninformative background.

Functionally, both the Decision module in Dispider and the ATS in Holmes-VAU serve as asynchronous gatekeepers similar to our Oscillatory Working Memory (OWM): they decouple continuous signal monitoring from expensive autoregressive inference to resolve the accuracy-cost trade-off. However, a key distinction lies in the detection mechanism. While the samplers in Dispider and Holmes-VAU typically rely on representation-based or supervised detectors that require dataset-specific training or historical data, our OWM approach operates training-free by maintaining a compact online state via oscillatory dynamics, eliminating the need for long-term historical buffering or offline calibration while achieving comparable or superior detection sensitivity.

B. Carrier Frequency Assignment

B.1. Frequency Allocation Strategy

We assign each PANN dimension $i \in \{0, 1, \dots, 526\}$ a distinct carrier frequency f_i to enable frequency-multiplexed oscillatory drives within the OWM lattice. The allocation follows a deterministic procedure ensuring compatibility with the discrete wave dynamics.

B.1.1. LINEAR FREQUENCY DISTRIBUTION

Given $C = 527$ PANN dimensions and a target frequency band $[f_{\min}, f_{\max}]$, carrier frequencies are uniformly distributed:

$$f_i = f_{\min} + (f_{\max} - f_{\min}) \times \frac{i}{C-1}, \quad i \in \{0, 1, \dots, 526\}. \quad (\text{B.1})$$

In our implementation, $f_{\min} = 51$ Hz and $f_{\max} = 1200$ Hz, yielding an average frequency spacing of $\Delta f \approx 2.18$ Hz.

B.1.2. FREQUENCY-TO-GRID MAPPING

Each frequency f_i is mapped to a spatial parcel $\Omega_i \subset \{0, \dots, G-1\} \times \{0, \dots, G-1\}$ on the $G \times G$ lattice, where $G = 64$.

Grid enumeration. All grid coordinates are enumerated in row-major order:

$$\mathcal{G} = \{(x, y) : x \in \{0, \dots, G-1\}, y \in \{0, \dots, G-1\}\}, \quad (\text{B.2})$$

where $|\mathcal{G}| = G^2 = 4096$.

Parcel allocation. The grid \mathcal{G} is partitioned into C contiguous parcels $\{\Omega_i\}_{i=0}^{C-1}$. To distribute the G^2 grid points as evenly as possible, we first compute:

$$n_{\text{base}} = \left\lfloor \frac{G^2}{C} \right\rfloor = 7, \quad n_{\text{rem}} = G^2 \bmod C = 407. \quad (\text{B.3})$$

The first n_{rem} parcels receive one additional grid point to accommodate the remainder:

$$|\Omega_i| = \begin{cases} n_{\text{base}} + 1 = 8 & \text{if } i < n_{\text{rem}}, \\ n_{\text{base}} = 7 & \text{if } i \geq n_{\text{rem}}. \end{cases} \quad (\text{B.4})$$

For each parcel i , the boundaries are determined by cumulative allocation:

$$\text{start}_i = \sum_{j=0}^{i-1} |\Omega_j|, \quad (\text{B.5})$$

$$\text{end}_i = \text{start}_i + |\Omega_i|. \quad (\text{B.6})$$

The parcel Ω_i then consists of grid positions with linear indices in $[\text{start}_i, \text{end}_i)$:

$$\Omega_i = \{(x, y) \in \mathcal{G} : \text{start}_i \leq x \cdot G + y < \text{end}_i\}. \quad (\text{B.7})$$

Note that $\sum_{i=0}^{C-1} |\Omega_i| < G^2$, with the remaining grid points left unassigned.

B.1.3. WAVE PROPAGATION SPEED COMPUTATION

For each parcel Ω_i , the local wave propagation speed c_i is computed from the target frequency f_i using the eigenfrequency relation from Theorem 2.1 Eq. (D.1) in Appendix D:

$$c_i = \frac{\tan(\pi f_i \Delta t) \sqrt{(1 + \Delta t k_p)(1 + \Delta t k_v)}}{\Delta t \sqrt{2}}, \quad (\text{B.8})$$

where $\Delta t = 0.01$ is the time step, $k_p = k_v = 10.0$ are damping coefficients, and the factor $\sqrt{2}$ accounts for the 2D isotropic spatial discretization with $dx = 1$.

Stability constraints. To ensure numerical stability, the computed wave propagation speed is clamped:

$$c_{\max} = \frac{dx}{\Delta t \sqrt{2}} \sqrt{(1 + \Delta t k_p)(1 + \Delta t k_v)} \approx 77.8, \quad (\text{B.9})$$

$$c_i \leftarrow \min(c_i, 0.9 \times c_{\max}) \approx 70.0, \quad (\text{B.10})$$

$$c_i \leftarrow \max(c_i, 0.1). \quad (\text{B.11})$$

The upper bound prevents CFL-like violations, while the lower bound ensures non-zero propagation for high-frequency modes.

Spatial wave propagation speed field. The wave propagation speed field $c(x, y)$ on the lattice is constructed by assigning c_i to all positions in parcel Ω_i :

$$c(x, y) = \sum_{i=0}^{526} c_i \cdot \mathcal{K}_{\Omega_i}(x, y), \quad (\text{B.12})$$

where $\mathcal{K}_{\Omega_i}(x, y)$ is the indicator function. This produces a spatially-varying field with approximately binary contrast between low-frequency (small c_i) and high-frequency (large c_i) parcels.

Nyquist Constraint. The lower bound $f_{\min} = 51$ Hz is chosen to strictly exceed the Nyquist frequency $f_N = 1/(2\Delta t) = 50$ Hz imposed by the discrete time step $\Delta t = 0.01$ s. Setting $f_{\min} = f_N$ would cause the drive signal $\sin(2\pi f_{\min} t)$ to evaluate to zero at all discrete timesteps $t = n\Delta t$ (since $\sin(\pi n) = 0$), and would create a singularity in the wave propagation speed formula (Eq. B.8) where $\tan(\pi f_{\min} \Delta t) = \tan(\pi/2) \rightarrow \infty$. The choice of 51 Hz maintains the original frequency spacing of approximately 2.18 Hz while ensuring well-defined dynamics.

B.2. Oscillatory Drive Signal Construction

At each time step t , the modulated excitation for PANN dimension i is:

$$S_i(x, y, t) = a_i(t) \sin(2\pi f_i t) \cdot \mathcal{K}_{\Omega_i}(x, y), \quad (\text{B.13})$$

where $a_i(t) \in [0, 1]$ is the PANN probability (sigmoid-activated logit) for event category i at time t . The total drive signal is the superposition:

$$S(x, y, t) = \sum_{i=0}^{526} S_i(x, y, t). \quad (\text{B.14})$$

Batch implementation. In the batch-aware implementation, time t is tracked independently for each sample $b \in \{1, \dots, B\}$:

$$t^{(b)} = \text{time_steps}^{(b)} \times \Delta t, \quad (\text{B.15})$$

where $\text{time_steps}^{(b)}$ is an integer counter maintained in `_batch_states[time_steps]`. The drive signal for batch b is:

$$S^{(b)}(x, y, t^{(b)}) = \sum_{i=0}^{526} a_i^{(b)}(t^{(b)}) \sin(2\pi f_i t^{(b)}) \cdot \mathcal{K}_{\Omega_i}(x, y). \quad (\text{B.16})$$

B.3. Nyquist Constraint and Phase Encoding

The discrete-time system with $\Delta t = 0.01$ s imposes a Nyquist frequency:

$$f_{\text{Nyquist}} = \frac{1}{2\Delta t} = 50 \text{ Hz}. \quad (\text{B.17})$$

This constrains the internal oscillatory dynamics of the pressure field $p(x, y, t)$ to frequencies below 50 Hz. FFT analysis (Appendix J) confirms dominant p -field frequencies lie within theta (4–8 Hz), alpha (8–12 Hz), beta (13–30 Hz), and low-gamma (30–50 Hz) bands.

High-frequency drive encoding. Drive frequencies $f_i > 50$ Hz do not produce corresponding oscillations in $p(x, y, t)$ due to the Nyquist constraint. Instead, these drives encode information through phase relationships. For integer time indices $n = t/\Delta t$, the drive signal phase:

$$\phi_i(n) = 2\pi f_i n \Delta t \pmod{2\pi} \tag{B.18}$$

varies uniquely for each f_i , enabling discrimination between 527 semantic categories via phase-based interference patterns and energy accumulation rates, despite internal dynamics remaining below 50 Hz.

B.4. Implementation Parameters

The frequency assignment is fully determined by the following fixed parameters:

- PANN encoder output dimension: $C = 527$
- Grid resolution: $G = 64$
- Time step: $\Delta t = 0.01$ s
- Spatial step: $dx = 1$
- Damping coefficients: $k_p = k_v = 10.0$
- Frequency range: $[f_{\min}, f_{\max}] = [50, 1200]$ Hz

These parameters remain constant across all experiments and datasets (XD-Violence, USoW), requiring no domain-specific calibration or retraining.

C. Detailed Algorithms

This section provides the complete algorithmic specifications for the OWM-based drift detection framework presented in Subsection 2.1. We present two key algorithms: the main drift detection pipeline and the adaptive threshold computation mechanism.

C.1. Main Saliency Detection Algorithm

Algorithm C.1 describes the complete workflow for detecting auditory pattern drift using the NAACA. The algorithm processes streaming audio input through several key stages: feature extraction via a pretrained encoder, oscillatory signal generation, OWM state updates, energy-based change detection, and persistence filtering.

The energy rate (first derivative) within the well-designed structure of the OWM system captures sudden transitions in the rate of energy change, which correspond to significant shifts in auditory patterns. The algorithm employs an adaptive threshold mechanism that automatically adjusts to the system’s baseline behavior and temporal trends, eliminating the need for manual threshold tuning.

To ensure robust detection and minimize false positives, the algorithm incorporates persistence filtering that requires a minimum proportion of recent detections before confirming a drift event. Additionally, a cooldown period prevents redundant detections of the same drift event.

C.2. Adaptive Threshold Computation

Algorithm C.2 details the online computation of adaptive thresholds for energy-based drift detection. The threshold adapts to both the statistical properties of recent observations and temporal trends in the data.

The algorithm maintains a sliding window of recent metric values and computes the threshold as $T = \mu + 2\sigma(1 + \alpha \cdot f_{\text{trend}})$, where μ and σ are the running mean and standard deviation, α is the trend adjustment factor, and f_{trend} quantifies the strength of temporal trends using linear regression.

During the initial bootstrap phase with insufficient data (fewer than 5 samples), the algorithm uses a simplified threshold computation to avoid instability. The trend factor captures whether the metric values are systematically increasing or

Algorithm C.1 OWM-based Auditory Pattern Drift Detection

Require: Audio stream $\{\mathbf{x}_t\}$, encoder $\text{Enc}(\cdot)$, OWM model with spatial operator \mathcal{A}
Require: Persistence duration $P = 3$, cooldown period $C = 3$

```

1: Initialize energy calculator
2: Initialize adaptive threshold  $T_{\text{adapt}}$  for energy metric
3: Initialize detection buffer  $\mathcal{D} = \emptyset$ , cooldown timer  $t_{\text{last}} = -1$ 
4: for each time step  $t$  do
5:    $\mathbf{p}_t \leftarrow \text{Enc}(\mathbf{x}_t)$   $\triangleright$  Extract event probabilities
6:   Generate oscillatory inputs:  $S_i(x, t) = a_i(t) \sin(\omega_i t) \mathbf{1}_{\Omega_i}(x)$ 
7:   Update OWM state with oscillatory drive signals
8:   Calculate current energy change rate from OWM dynamics  $\triangleright$  Energy change rate
9:    $T_{\text{adapt}} \leftarrow \mu + 2\sigma(1 + \alpha \cdot \text{trend})$   $\triangleright$  Update adaptive threshold
10:   $d_{\text{candidate}} \leftarrow \mathbf{1}(\text{energy change rate} > T_{\text{adapt}})$   $\triangleright$  Energy-based detection
11:  Add  $d_{\text{candidate}}$  to detection buffer  $\mathcal{D}$   $\triangleright$  Persistence filtering
12:  if  $|\mathcal{D}| \geq P$  then
13:     $r_{\text{persist}} \leftarrow \frac{1}{P} \sum_{i=0}^{P-1} \mathcal{D}[-P+i]$   $\triangleright$  Persistence ratio
14:    if  $r_{\text{persist}} \geq 0.5$  AND  $t - t_{\text{last}} > C$  then
15:      output Drift detected at time  $t$ 
16:       $t_{\text{last}} \leftarrow t$ , clear  $\mathcal{D}$   $\triangleright$  Reset detection state
17:    end if
18:  end if
19: end for

```

Algorithm C.2 Online Adaptive Threshold Calculation

Require: Window size $W = 20$, trend adjustment factor $\alpha = 0.2$

```

1: Initialize value buffer  $\mathcal{V} = \emptyset$  with maximum size  $W$ 
2: function AdaptiveThreshold.update( $v_{\text{new}}$ )
3: Add  $v_{\text{new}}$  to buffer  $\mathcal{V}$ 
4: if  $|\mathcal{V}| < 5$  then
5:    $\mu \leftarrow \frac{1}{|\mathcal{V}|} \sum_i \mathcal{V}_i$   $\triangleright$  Insufficient data for robust statistics
6:   if  $|\mathcal{V}| > 1$  then
7:      $\sigma \leftarrow \sqrt{\frac{1}{|\mathcal{V}|-1} \sum_i (\mathcal{V}_i - \mu)^2}$ 
8:   else
9:      $\sigma \leftarrow 0.1$ 
10:  end if
11:  return  $\mu + 1.5\sigma$   $\triangleright$  Simple threshold for bootstrap
12: end if  $\triangleright$  Compute baseline statistics
13:  $\mu \leftarrow \frac{1}{|\mathcal{V}|} \sum_i \mathcal{V}_i$   $\triangleright$  Running mean
14:  $\sigma \leftarrow \sqrt{\frac{1}{|\mathcal{V}|-1} \sum_i (\mathcal{V}_i - \mu)^2}$   $\triangleright$  Running std
15: if  $|\mathcal{V}| \geq 3$  then
16:    $\mathbf{x} \leftarrow [0, 1, \dots, |\mathcal{V}| - 1]$   $\triangleright$  Time indices
17:    $\text{slope} \leftarrow \frac{\sum_i (x_i - \bar{x})(\mathcal{V}_i - \mu)}{\sum_i (x_i - \bar{x})^2}$   $\triangleright$  Linear regression slope
18:    $f_{\text{trend}} \leftarrow \frac{|\text{slope}|}{\sigma + 10^{-8}}$   $\triangleright$  Normalized trend strength
19: else
20:    $f_{\text{trend}} \leftarrow 0$ 
21: end if
22:  $T_{\text{adapt}} \leftarrow \mu + 2\sigma(1 + \alpha \cdot f_{\text{trend}})$   $\triangleright$  Adaptive threshold with trend adjustment
23: return  $T_{\text{adapt}}$ 

```

decreasing, allowing the threshold to adapt accordingly. This prevents false negatives during periods of natural system evolution while maintaining sensitivity to abrupt changes.

The adaptive nature of this threshold computation is crucial for handling diverse auditory environments with varying baseline activity levels, ensuring consistent detection performance across different acoustic contexts without requiring environment-specific calibration.

D. Proof for Theorem 2.1

D.1. Eigenvalue Analysis and Local Resonance Frequency

The OWM system is governed by the discrete formulation in Eqs. 3. For a local grid point (x, y) , define the state vector $\mathbf{h} = [p, v_x, v_y]^T$. The update can be written as

$$\mathbf{h}(t + \Delta t) = \mathbf{A}\mathbf{h}(t) + S(x, y, t), \quad (\text{D.1})$$

where

$$\mathbf{A} = \mathbf{M}_{\text{Damp}}^{-1} \mathbf{M}_{\text{Velocity}}, \quad (\text{D.2})$$

$$\mathbf{M}_{\text{Velocity}} = \begin{bmatrix} 1 & -c^2 \Delta t \frac{\partial}{\partial x} & -c^2 \Delta t \frac{\partial}{\partial y} \\ -\Delta t \frac{\partial}{\partial x} & 1 & 0 \\ -\Delta t \frac{\partial}{\partial y} & 0 & 1 \end{bmatrix}, \quad (\text{D.3})$$

$$\mathbf{M}_{\text{Damp}} = \begin{bmatrix} 1 + \Delta t k^p(x, y) & 0 & 0 \\ 0 & 1 + \Delta t k^v(x, y) & 0 \\ 0 & 0 & 1 + \Delta t k^v(x, y) \end{bmatrix}. \quad (\text{D.4})$$

In Fourier space, $\partial/\partial x \rightarrow i\xi_x$ and $\partial/\partial y \rightarrow i\xi_y$. The corresponding system matrix becomes

$$\mathbf{A} = \mathbf{M}_{\text{Damp}}^{-1} \mathbf{M}_{\text{Velocity}} = \begin{bmatrix} \frac{1}{1 + \Delta t k_{i,j}^p} & -\frac{c_{i,j}^2 \Delta t i \xi_x}{1 + \Delta t k_{i,j}^p} & -\frac{c_{i,j}^2 \Delta t i \xi_y}{1 + \Delta t k_{i,j}^p} \\ -\frac{\Delta t i \xi_x}{1 + \Delta t k_{i,j}^v} & \frac{1}{1 + \Delta t k_{i,j}^v} & 0 \\ -\frac{\Delta t i \xi_y}{1 + \Delta t k_{i,j}^v} & 0 & \frac{1}{1 + \Delta t k_{i,j}^v} \end{bmatrix}. \quad (\text{D.5})$$

The characteristic equation $\det(\mathbf{A} - \lambda \mathbf{I}) = 0$ expands to

$$\left(\frac{1}{1 + \Delta t k_{i,j}^v} - \lambda \right)^2 \left[\left(\frac{1}{1 + \Delta t k_{i,j}^p} - \lambda \right) \left(\frac{1}{1 + \Delta t k_{i,j}^v} - \lambda \right) + \frac{c_{i,j}^2 \Delta t^2 (\xi_x^2 + \xi_y^2)}{(1 + \Delta t k_{i,j}^p)(1 + \Delta t k_{i,j}^v)} \right] = 0. \quad (\text{D.6})$$

This yields one real eigenvalue $\lambda_1 = 1/(1 + \Delta t k_{i,j}^v)$ and two eigenvalues satisfying

$$\left(\frac{1}{1 + \Delta t k_{i,j}^p} - \lambda \right) \left(\frac{1}{1 + \Delta t k_{i,j}^v} - \lambda \right) + \frac{c_{i,j}^2 \Delta t^2 (\xi_x^2 + \xi_y^2)}{(1 + \Delta t k_{i,j}^p)(1 + \Delta t k_{i,j}^v)} = 0. \quad (\text{D.7})$$

For the parameter regime used ($\Delta t k_p = \Delta t k_v = 0.1 \ll 1$ and typical spatial wavelengths), the discriminant of Eq. D.7 satisfies $\Delta < 0$, ensuring complex conjugate eigenvalues. Under the balanced damping condition $k_p = k_v$, the real parts cancel exactly, yielding:

$$\lambda_{2,3} = \frac{1}{2} \left[\frac{1}{1 + \Delta t k_{i,j}^p} + \frac{1}{1 + \Delta t k_{i,j}^v} \right] \pm i \frac{c_{i,j} \Delta t \sqrt{\xi_x^2 + \xi_y^2}}{\sqrt{(1 + \Delta t k_{i,j}^p)(1 + \Delta t k_{i,j}^v)}}. \quad (\text{D.8})$$

Let θ be the phase angle of $\lambda_{2,3}$:

$$\theta = \tan^{-1} \left(\frac{\text{Im}(\lambda)}{\text{Re}(\lambda)} \right) = \tan^{-1} \left(\frac{2c_{i,j} \Delta t \sqrt{\xi_x^2 + \xi_y^2}}{\sqrt{(1 + \Delta t k_{i,j}^p)(1 + \Delta t k_{i,j}^v)} \left(\frac{1}{1 + \Delta t k_{i,j}^p} + \frac{1}{1 + \Delta t k_{i,j}^v} \right)} \right). \quad (\text{D.9})$$

In discrete time, the oscillation frequency associated with phase θ is $f = \theta/(\pi\Delta t)$, hence the local characteristic frequency is

$$f_{i,j} = \frac{1}{\pi\Delta t} \tan^{-1} \left(\frac{2c_{i,j}\Delta t \sqrt{\xi_x^2 + \xi_y^2}}{\sqrt{(1 + \Delta t k_{i,j}^p)(1 + \Delta t k_{i,j}^v) \left(\frac{1}{1 + \Delta t k_{i,j}^p} + \frac{1}{1 + \Delta t k_{i,j}^v} \right)}} \right). \quad (\text{D.10})$$

We define the local resonance frequency of the discrete system by $\omega_{\text{res}}(x, y) = 2\pi f_{i,j}$.

D.2. Frequency-Domain Response and Phase Delay

Consider harmonic excitation at angular frequency ω . In the frequency domain, the continuous system in Eqs. (2) satisfies

$$[i\omega + k^p(x, y)]p(x, y, \omega) = -c^2(x, y)\nabla \cdot \mathbf{v}(x, y, \omega) + S_0(x, y), \quad (\text{D.11})$$

$$[i\omega + k^v(x, y)]\mathbf{v}(x, y, \omega) = -\nabla p(x, y, \omega). \quad (\text{D.12})$$

Eliminating \mathbf{v} gives $[i\omega + k^p][i\omega + k^v]p - c^2\nabla^2 p = [i\omega + k^v]S_0$.

The resonance frequency $\omega_{\text{res}}(x, y)$ is the natural angular frequency of the discrete system from the previous subsection. Using $\omega_{\text{res}} = \theta/\Delta t$ and the small Δt approximation yields

$$\omega_{\text{res}}(x, y) \approx c(x, y)\sqrt{\xi_x^2 + \xi_y^2}. \quad (\text{D.13})$$

For slowly varying parameters, we use the local approximation $-c^2\nabla^2 \rightarrow \omega_{\text{res}}^2(x, y)$, leading to the localized frequency-domain relation

$$\left[(\omega_{\text{res}}^2 - k^p k^v - \omega^2) + i\omega(k^p + k^v) \right] p \approx (i\omega + k^v)S_0. \quad (\text{D.14})$$

Define the effective local resonance by $\Omega_{\text{res}}^2(x, y) \triangleq \omega_{\text{res}}^2(x, y) - k^p(x, y)k^v(x, y)$. Then the complex denominator governing the response is

$$\mathcal{D}(\omega) = (\Omega_{\text{res}}^2 - \omega^2) + i\omega(k^p + k^v). \quad (\text{D.15})$$

Hence the amplitude response satisfies $A(x, y, \omega) = |i\omega + k^v|/|\mathcal{D}(\omega)|$.

The phase of p is given by the argument of $\mathcal{D}(\omega)$. Writing $\phi(\omega; x, y) = -\arg \mathcal{D}(\omega)$, one obtains

$$\phi(\omega; x, y) = -\arctan \left(\frac{\omega(k^p + k^v)}{\Omega_{\text{res}}^2 - \omega^2} \right). \quad (\text{D.16})$$

Near resonance, linearize $\Omega_{\text{res}}^2 - \omega^2 \approx -2\Omega_{\text{res}}(\omega - \Omega_{\text{res}})$ to obtain

$$\phi(\omega; x, y) \approx \arctan \left(\frac{\omega(k^p + k^v)}{2\Omega_{\text{res}}(\omega - \Omega_{\text{res}})} \right). \quad (\text{D.17})$$

Up to a slowly varying prefactor (absorbed by the proportionality in Eq. (5)), this yields the temporal phase delay. Using the identity $\arctan(-x) = -\arctan(x)$ and absorbing the sign into the definition of positive phase delay (where lag is defined as positive), we obtain:

$$\tau(\omega; x, y) = \frac{\phi(\omega; x, y)}{\omega} \propto \arctan \left(\frac{\omega - \omega_{\text{res}}(x, y)}{k^p(x, y) + k^v(x, y)} \right), \quad (\text{D.18})$$

which matches the form in Eq. (5).

Finally, by linearity and superposition, for multi-frequency excitation $S(x, y, t) = \sum_i S_i(x, y, t)$ with narrowband components at ω_i and under the quasi-steady assumption (envelope variation slower than $(k^p + k^v)^{-1}$), each component is locally delayed, giving

$$p(x, y, t) \approx \sum_i A_i(x, y) S_i(x, y, t - \tau_i(x, y)), \quad (\text{D.19})$$

which establishes Eq. (4).

D.3. Application to OWM and Generality

In OWM (Section 3 and Appendix B), $c(x, y)$ is designed to vary spatially, which makes $\omega_{\text{res}}(x, y) = 2\pi f_{i,j}$ spatially varying through Eq. (D.10). For a component at frequency ω_i , Eq. (D.18) gives a location-dependent delay $\tau_i(x, y)$, and superposition yields the delayed-sum form in Eq. (4).

The proof relies on (i) linear operation so that superposition holds, (ii) the quasi-steady condition on the excitation envelope relative to $(k^p + k^v)^{-1}$, and (iii) slowly varying spatial parameters so that the local approximation in Eq. (D.14) is valid. Under these conditions, different input frequencies induce different local phase delays determined by the mismatch between ω_i and $\omega_{\text{res}}(x, y)$.

E. Proof for Theorem 2.2

Proof. We prove the equivalence by transforming the system step by step. We begin by differentiating the pressure equation from Eqs. 2 with respect to time:

$$\frac{\partial^2 p}{\partial t^2} + k^p \frac{\partial p}{\partial t} = -\frac{\partial}{\partial t} [\nabla \cdot (c^2(x, y)\mathbf{v})] + \frac{\partial S}{\partial t}. \quad (\text{E.1})$$

Since $c^2(x, y)$ is time-independent, we apply the product rule for the divergence of a time-dependent vector field with spatially-varying coefficients:

$$\frac{\partial}{\partial t} [\nabla \cdot (c^2\mathbf{v})] = \nabla \cdot \left(c^2 \frac{\partial \mathbf{v}}{\partial t} \right) + \frac{\partial \mathbf{v}}{\partial t} \cdot \nabla c^2. \quad (\text{E.2})$$

To evaluate the velocity derivatives, we use the velocity equations from Eqs. 2:

$$\frac{\partial \mathbf{v}}{\partial t} = -k^v \mathbf{v} - \nabla p. \quad (\text{E.3})$$

Substituting Eq. E.3 into the divergence term in Eq. E.2:

$$\begin{aligned} \nabla \cdot \left(c^2 \frac{\partial \mathbf{v}}{\partial t} \right) &= \nabla \cdot [c^2(-k^v \mathbf{v} - \nabla p)] \\ &= -k^v \nabla \cdot (c^2 \mathbf{v}) - \nabla \cdot (c^2 \nabla p). \end{aligned} \quad (\text{E.4})$$

For the gradient coupling term in Eq. E.2:

$$\frac{\partial \mathbf{v}}{\partial t} \cdot \nabla c^2 = (-k^v \mathbf{v} - \nabla p) \cdot \nabla c^2 = -k^v \mathbf{v} \cdot \nabla c^2 - \nabla p \cdot \nabla c^2. \quad (\text{E.5})$$

Combining Eqs. E.4 and E.5 into Eq. E.2:

$$\frac{\partial}{\partial t} [\nabla \cdot (c^2\mathbf{v})] = -k^v \nabla \cdot (c^2\mathbf{v}) - \nabla \cdot (c^2 \nabla p) - k^v \mathbf{v} \cdot \nabla c^2 - \nabla p \cdot \nabla c^2. \quad (\text{E.6})$$

We now eliminate the velocity divergence using the pressure equation from Eqs. 2:

$$\nabla \cdot (c^2\mathbf{v}) = -\frac{\partial p}{\partial t} - k^p p + S. \quad (\text{E.7})$$

Inserting Eq. E.7 into Eq. E.6:

$$\begin{aligned}
 \frac{\partial}{\partial t}[\nabla \cdot (c^2 \mathbf{v})] &= -k^v \left(-\frac{\partial p}{\partial t} - k^p p + S \right) - \nabla \cdot (c^2 \nabla p) \\
 &\quad - k^v \mathbf{v} \cdot \nabla c^2 - \nabla p \cdot \nabla c^2 \\
 &= k^v \frac{\partial p}{\partial t} + k^v k^p p - k^v S - \nabla \cdot (c^2 \nabla p) \\
 &\quad - k^v \mathbf{v} \cdot \nabla c^2 - \nabla p \cdot \nabla c^2.
 \end{aligned} \tag{E.8}$$

Substituting Eq. E.8 back into Eq. E.1:

$$\begin{aligned}
 \frac{\partial^2 p}{\partial t^2} + k^p \frac{\partial p}{\partial t} &= - \left(k^v \frac{\partial p}{\partial t} + k^v k^p p - k^v S - \nabla \cdot (c^2 \nabla p) \right. \\
 &\quad \left. - k^v \mathbf{v} \cdot \nabla c^2 - \nabla p \cdot \nabla c^2 \right) + \frac{\partial S}{\partial t}.
 \end{aligned} \tag{E.9}$$

Finally, rearranging terms leads to:

$$\frac{\partial^2 p}{\partial t^2} + (k^p + k^v) \frac{\partial p}{\partial t} + k^v k^p p = \nabla \cdot (c^2 \nabla p) + (k^v \mathbf{v} + \nabla p) \cdot \nabla c^2 + \left(k^v + \frac{\partial}{\partial t} \right) S, \tag{E.10}$$

which matches Eq. 8 with effective damping coefficient $\gamma = k^p + k^v$, restoring force coefficient $\mu = k^v k^p$, and spatial coupling term $(k^v \mathbf{v} + \nabla p) \cdot \nabla c^2$ arising from the spatially-varying wave speed. \square

F. Proof for Theorem 2.3

Proof. We begin by computing the time derivative of energy:

$$\frac{dE}{dt} = \iint \left[p \frac{\partial p}{\partial t} + v_x \frac{\partial v_x}{\partial t} + v_y \frac{\partial v_y}{\partial t} \right] dx dy. \tag{F.1}$$

Substituting the system equations from Eqs. 2 into Eq. F.1, we obtain:

$$\frac{dE}{dt} = \iint \left[p \left(-k^p p - c^2 \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) + S \right) + v_x \left(-k^v v_x - \frac{\partial p}{\partial x} \right) + v_y \left(-k^v v_y - \frac{\partial p}{\partial y} \right) \right] dx dy. \tag{F.2}$$

Expanding the terms in Eq. F.2 yields:

$$\frac{dE}{dt} = \iint \left[-k^p p^2 - c^2 p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) + pS - k^v v_x^2 - v_x \frac{\partial p}{\partial x} - k^v v_y^2 - v_y \frac{\partial p}{\partial y} \right] dx dy. \tag{F.3}$$

We can group the terms in Eq. F.3 to separate dissipation, coupling, and source contributions:

$$\frac{dE}{dt} = \iint \left[-k^p p^2 - k^v (v_x^2 + v_y^2) \right] dx dy + \iint \left[-c^2 p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) - v_x \frac{\partial p}{\partial x} - v_y \frac{\partial p}{\partial y} \right] dx dy + \iint pS dx dy. \tag{F.4}$$

To handle the coupling terms in Eq. F.4, we apply integration by parts. For example:

$$\iint v_x \frac{\partial p}{\partial x} dx dy = v_x p \Big|_{x \text{ boundaries}} - \iint p \frac{\partial v_x}{\partial x} dx dy. \tag{F.5}$$

Since periodic boundary conditions imply the boundary term vanishes, we have:

$$\iint v_x \frac{\partial p}{\partial x} dx dy = - \iint p \frac{\partial v_x}{\partial x} dx dy. \tag{F.6}$$

Similarly, for the y -component:

$$\iint v_y \frac{\partial p}{\partial y} dx dy = - \iint p \frac{\partial v_y}{\partial y} dx dy. \quad (\text{F.7})$$

Using Eqs. F.6 and F.7, the coupling terms in Eq. F.4 reduce to:

$$\begin{aligned} & -c^2 p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) - v_x \frac{\partial p}{\partial x} - v_y \frac{\partial p}{\partial y} \\ & = -c^2 p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) + p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) \\ & = -(c^2 - 1) p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right). \end{aligned} \quad (\text{F.8})$$

Combining all results, we obtain the final expression:

$$\frac{dE}{dt} = - \iint [k^p p^2 + k^v (v_x^2 + v_y^2)] dx dy - \iint (c^2 - 1) p \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} \right) dx dy + \iint p S dx dy, \quad (\text{F.9})$$

which establishes the claim. \square

G. Proofs for wave propagation speed Optimization

G.1. Interface Reflection and Contrast Effects

Lemma G.1 (Interface Reflection Coefficient). *Consider a plane wave incident on an interface between two constant-speed media with speeds c_1 and c_2 . With uniform density $\rho = 1$ so that impedance $Z = \rho c = c$, the pressure reflection coefficient is*

$$R = \frac{c_2 - c_1}{c_1 + c_2}. \quad (\text{G.1})$$

The energy reflection ratio is $E_{\text{reflected}}/E_{\text{incident}} = |R|^2$. Under a fixed speed-sum constraint $c_1 + c_2 = C$, $|R|^2$ is maximized by extreme (binary) contrasts $(c_1, c_2) \in \{(0, C), (C, 0)\}$, while $c_1 = c_2 = C/2$ yields $R = 0$.

Proof. With $\rho = 1$, the medium impedances are $Z_i = c_i$, $i = 1, 2$. Continuity of pressure and normal velocity (using $v = p/Z$) at the interface yields

$$p_1 + p_1^r = p_2^t, \quad \frac{p_1}{Z_1} - \frac{p_1^r}{Z_1} = \frac{p_2^t}{Z_2}, \quad (\text{G.2})$$

where subscripts denote incident (no subscript), reflected (r), and transmitted (t) components. Solving these equations gives

$$p_2^t = \frac{2Z_2}{Z_1 + Z_2} p_1, \quad \frac{p_1^r}{p_1} = \frac{Z_2 - Z_1}{Z_1 + Z_2} = \frac{c_2 - c_1}{c_1 + c_2} = R. \quad (\text{G.3})$$

The reflected-to-incident energy ratio is therefore

$$\frac{E_{\text{reflected}}}{E_{\text{incident}}} = |R|^2 = \left(\frac{c_2 - c_1}{c_1 + c_2} \right)^2. \quad (\text{G.4})$$

If $c_1 + c_2 = C$, maximizing $|R|^2 = \left(\frac{c_2 - c_1}{C} \right)^2$ is equivalent to maximizing $|c_2 - c_1|$. The maximum occurs at the binary extremes $(c_1, c_2) = (0, C)$ or $(C, 0)$, for which $|R| \rightarrow 1$. Any interior choice yields a smaller R . For example, $c_1 = c_2 = C/2$ gives $R = 0$, while $c_1 = 0.9C$, $c_2 = 0.1C$ gives $|R| = 0.8 < 1$. \square

G.2. Modal Coupling Framework

Lemma G.2 (Modal Coupling Derivation and Selection Rule). *In the periodic domain $[0, L_x] \times [0, L_y]$ with orthonormal Fourier basis*

$$\phi_{m,n}(x, y) = \frac{1}{\sqrt{A}} \exp\left(i \frac{2\pi m x}{L_x} + i \frac{2\pi n y}{L_y}\right), \quad A = L_x L_y, \quad (\text{G.5})$$

and wavenumbers

$$k_{m,n}^2 = \left(\frac{2\pi m}{L_x}\right)^2 + \left(\frac{2\pi n}{L_y}\right)^2, \quad (\text{G.6})$$

write $p(x, y, t) = \sum_{m,n} a_{m,n}(t) \phi_{m,n}(x, y)$ and $c^2(x, y) = c_0^2 + \delta c^2(x, y)$. Here, $a_{m,n}(t)$ are the time-dependent modal coefficients. Neglecting damping, the modal system

$$\ddot{a}_{m,n} + \omega_{m,n}^2 a_{m,n} = \sum_{m',n'} C_{(m,n),(m',n')} a_{m',n'} + S_{m,n}, \quad (\text{G.7})$$

has $\omega_{m,n}^2 = c_0^2 k_{m,n}^2$ and coupling

$$C_{(m,n),(m',n')} = k_{m',n'}^2 \langle \delta c^2, \phi_{m',n'}, \phi_{m,n} \rangle. \quad (\text{G.8})$$

For striped $\delta c^2(x, y) = f(y)$,

$$\langle \delta c^2, \phi_{m',n'}, \phi_{m,n} \rangle = \delta_{m,m'} \hat{V}_{n-n'}, \quad (\text{G.9})$$

i.e., coupling is block-diagonal in m and depends on the y -Fourier coefficients \hat{V}_q of f , where

$$\hat{V}_q = \frac{1}{L_y} \int_0^{L_y} f(y) e^{i \frac{2\pi q y}{L_y}} dy. \quad (\text{G.10})$$

Proof. We begin by establishing the modal expansion and deriving the coupling matrix through systematic projection onto Fourier basis functions.

On the periodic domain $[0, L_x] \times [0, L_y]$, we employ the orthonormal Fourier basis defined in Eq. (G.5), where $A = L_x L_y$ is the domain area. These functions satisfy $\nabla^2 \phi_{m,n} = -k_{m,n}^2 \phi_{m,n}$.

Starting from the wave equation without damping:

$$\frac{\partial^2 p}{\partial t^2} = c^2(x, y) \nabla^2 p + S(x, y, t), \quad c^2(x, y) = c_0^2 + \delta c^2(x, y), \quad (\text{G.11})$$

we expand $p(x, y, t) = \sum_{m,n} a_{m,n}(t) \phi_{m,n}(x, y)$ to obtain:

$$\begin{aligned} \sum_{m,n} \ddot{a}_{m,n}(t) \phi_{m,n}(x, y) &= (c_0^2 + \delta c^2(x, y)) \nabla^2 \left(\sum_{m,n} a_{m,n}(t) \phi_{m,n}(x, y) \right) + S(x, y, t) \\ &= \sum_{m,n} a_{m,n}(t) (c_0^2 + \delta c^2(x, y)) \nabla^2 \phi_{m,n}(x, y) + S(x, y, t). \end{aligned} \quad (\text{G.12})$$

Substituting $\nabla^2 \phi_{m,n} = -k_{m,n}^2 \phi_{m,n}$ into Eq. (G.12) yields:

$$\sum_{m,n} \ddot{a}_{m,n}(t) \phi_{m,n} = - \sum_{m,n} a_{m,n}(t) k_{m,n}^2 (c_0^2 + \delta c^2) \phi_{m,n} + S(x, y, t). \quad (\text{G.13})$$

To project this equation onto individual modes, we define the normalized inner product on $[0, L_x] \times [0, L_y]$ by:

$$\langle f, g \rangle := \frac{1}{A} \int_0^{L_x} \int_0^{L_y} f(x, y) g^*(x, y) dx dy, \quad A = L_x L_y, \quad (\text{G.14})$$

so that the Fourier modes are orthonormal: $\langle \phi_{m,n}, \phi_{\mu,\nu} \rangle = \delta_{m,\mu} \delta_{n,\nu}$, where δ denotes the Kronecker delta function. We also use the trilinear shorthand:

$$\langle \delta c^2, \phi_{m,n}, \phi_{\mu,\nu} \rangle := \frac{1}{A} \int_0^{L_x} \int_0^{L_y} \delta c^2(x, y) \phi_{m,n}^*(x, y) \phi_{\mu,\nu}(x, y) dx dy. \quad (\text{G.15})$$

Multiplying Eq. (G.13) by $\phi_{\mu,\nu}(x, y)$ and integrating over the domain, we analyze each term separately. For the left-hand side:

$$\left\langle \sum_{m,n} \ddot{a}_{m,n} \phi_{m,n}, \phi_{\mu,\nu} \right\rangle = \sum_{m,n} \ddot{a}_{m,n} \langle \phi_{m,n}, \phi_{\mu,\nu} \rangle = \sum_{m,n} \ddot{a}_{m,n} \delta_{m,\mu} \delta_{n,\nu} = \ddot{a}_{\mu,\nu}(t). \quad (\text{G.16})$$

For the reference-speed part:

$$\begin{aligned} - \left\langle \sum_{m,n} a_{m,n} k_{m,n}^2 c_0^2 \phi_{m,n}, \phi_{\mu,\nu} \right\rangle &= -c_0^2 \sum_{m,n} a_{m,n} k_{m,n}^2 \langle \phi_{m,n}, \phi_{\mu,\nu} \rangle \\ &= -c_0^2 k_{\mu,\nu}^2 a_{\mu,\nu}(t). \end{aligned} \quad (\text{G.17})$$

For the perturbation (coupling) part:

$$- \left\langle \sum_{m,n} a_{m,n} k_{m,n}^2 \delta c^2 \phi_{m,n}, \phi_{\mu,\nu} \right\rangle = - \sum_{m,n} a_{m,n} k_{m,n}^2 \langle \delta c^2, \phi_{m,n}, \phi_{\mu,\nu} \rangle. \quad (\text{G.18})$$

For the source part:

$$\langle S(x, y, t), \phi_{\mu,\nu} \rangle =: S_{\mu,\nu}(t). \quad (\text{G.19})$$

Collecting terms from Eqs. (G.16)–(G.19) gives:

$$\ddot{a}_{\mu,\nu}(t) + c_0^2 k_{\mu,\nu}^2 a_{\mu,\nu}(t) = - \sum_{m,n} k_{m,n}^2 \langle \delta c^2, \phi_{m,n}, \phi_{\mu,\nu} \rangle a_{m,n}(t) + S_{\mu,\nu}(t). \quad (\text{G.20})$$

Defining $\omega_{m,n}^2 := c_0^2 k_{m,n}^2$ and the coupling matrix:

$$C_{(\mu,\nu),(m,n)} := k_{m,n}^2 \langle \delta c^2, \phi_{m,n}, \phi_{\mu,\nu} \rangle, \quad (\text{G.21})$$

and relabeling $(\mu, \nu) \rightarrow (m, n)$, $(m, n) \rightarrow (m', n')$ in the sum, we obtain the modal system in Eq. (G.7).

For the special case of striped distributions where $\delta c^2(x, y) = f(y)$, we can evaluate the coupling matrix explicitly using Eq. (G.15):

$$\begin{aligned} \langle \delta c^2, \phi_{m',n'}, \phi_{m,n} \rangle &= \frac{1}{A} \iint f(y) e^{i \frac{2\pi(m-m')x}{L_x}} e^{i \frac{2\pi(n-n')y}{L_y}} dx dy \\ &= \frac{1}{A} \left[\int_0^{L_x} e^{i \frac{2\pi(m-m')x}{L_x}} dx \right] \left[\int_0^{L_y} f(y) e^{i \frac{2\pi(n-n')y}{L_y}} dy \right]. \end{aligned} \quad (\text{G.22})$$

The x -integral in Eq. (G.22) evaluates to $L_x \delta_{m,m'}$. Defining the Fourier coefficients:

$$\hat{V}_q = \frac{1}{L_y} \int_0^{L_y} f(y) e^{i \frac{2\pi q y}{L_y}} dy, \quad q = n - n', \quad (\text{G.23})$$

the y -integral becomes $L_y \hat{V}_{n-n'}$. Therefore:

$$\langle \delta c^2, \phi_{m',n'}, \phi_{m,n} \rangle = \delta_{m,m'} \hat{V}_{n-n'}. \quad (\text{G.24})$$

Substituting Eq. (G.24) into Eq. (G.21), the coupling matrix under striped distributions becomes block-diagonal in m :

$$C_{(m,n),(m',n')} = k_{m',n'}^2 \delta_{m,m'} \hat{V}_{n-n'}. \quad (\text{G.25})$$

□

G.3. Fourier Coefficient Maximization

Lemma G.3 (Square-Wave Fourier Optimality). *For a periodic function $f(y)$ with period L and pointwise constraint $|f(y)| \leq A$, the magnitude of the q -th Fourier coefficient*

$$\hat{f}_q = \frac{1}{L} \int_0^L f(y) e^{i2\pi qy/L} dy \quad (\text{G.26})$$

is maximized by the square wave

$$f(y) = A \cdot \text{sgn} \left[\cos \left(\frac{2\pi qy}{L} \right) \right], \quad (\text{G.27})$$

achieving $|\hat{f}_q| = 2A/(\pi q)$ for odd q .

Proof. By the Cauchy-Schwarz inequality,

$$|\hat{f}_q| \leq \frac{1}{L} \int_0^L |f(y)| dy \leq \frac{1}{L} \int_0^L A dy = A. \quad (\text{G.28})$$

Equality holds when $f(y)$ is phase-aligned with $e^{i2\pi qy/L}$, i.e., when

$$f(y) = A \cdot \text{sgn} \left[\cos \left(\frac{2\pi qy}{L} \right) \right]. \quad (\text{G.29})$$

For this square wave, direct computation gives

$$\hat{f}_q = \frac{1}{L} \int_0^L A \cdot \text{sgn} \left[\cos \left(\frac{2\pi qy}{L} \right) \right] e^{i2\pi qy/L} dy = \frac{2A}{\pi q} \quad (\text{G.30})$$

for odd q , confirming optimality.

Alternatively, this result follows from the convex optimization perspective: the set of functions satisfying $|f(y)| \leq A$ and fixed mean is convex, and the objective $|\hat{f}_q|$ is a linear functional. By the extreme point theorem, the maximum occurs at an extreme point of the feasible set, which corresponds to $f(y) \in \{-A, A\}$ almost everywhere. The phase-alignment condition then determines the switching pattern, yielding the square wave in Eq. (G.27). \square

G.4. Proof of Theorem 2.4

We now prove that the striped square-wave perturbation maximizes drift detection sensitivity through optimal spatial organization.

Proof. The proof establishes optimality through three complementary perspectives: geometric structure, resonant coupling, and spectral amplitude.

Part I: Geometric constraint from striped structure. From Lemma G.2, when the speed perturbation takes the form $\delta c^2(x, y) = f(y)$ (independent of x), the coupling matrix satisfies

$$C_{(m,n),(m',n')} = k_{m',n'}^2 \delta_{m,m'} \hat{V}_{n-n'}, \quad (\text{G.31})$$

where \hat{V}_q are the Fourier coefficients of $f(y)$ defined by

$$\hat{V}_q = \frac{1}{L_y} \int_0^{L_y} f(y) e^{i\frac{2\pi qy}{L_y}} dy. \quad (\text{G.32})$$

The Kronecker delta $\delta_{m,m'}$ in Eq. (G.31) reveals that the striped structure decouples modes with different x -direction wavenumbers: modes (m, n) and (m', n') with $m \neq m'$ do not couple. Within each m -block, coupling between modes (m, n) and (m, n') depends solely on the Fourier coefficient $\hat{V}_{n-n'}$ of the y -direction profile $f(y)$.

This geometric decoupling converts the multi-dimensional optimization problem over $c(x, y)$ into independent one-dimensional problems along y for each fixed m . This reduction reveals that unidirectional (striped) variation is structurally optimal for maximizing targeted modal coupling.

Part II: Resonant period selection via Fourier harmonics. To maximize coupling between modes (m, n) and $(m, n+q_0)$, we examine

$$|C_{(m,n),(m,n+q_0)}| = k_{m,n+q_0}^2 |\hat{V}_{q_0}|. \quad (\text{G.33})$$

Since $k_{m,n+q_0}^2$ is fixed by mode indices, maximizing coupling reduces to maximizing $|\hat{V}_{q_0}|$.

Consider a periodic function $f(y)$ on $[0, L_y]$ with period T (where T divides L_y evenly). The periodicity condition $f(y) = f(y+T)$ implies that the Fourier series Eq. (G.32) has $\hat{V}_q = 0$ unless q is an integer multiple of L_y/T . Specifically, the Fourier coefficients are nonzero only at

$$q = n \cdot \frac{L_y}{T}, \quad n = 0, \pm 1, \pm 2, \dots \quad (\text{G.34})$$

For the fundamental Fourier component (largest amplitude in a typical periodic function) to occur at the target mode index q_0 , we require

$$q_0 = 1 \cdot \frac{L_y}{T} \Rightarrow T = \frac{L_y}{q_0}. \quad (\text{G.35})$$

This period selection has a clear physical interpretation. In the modal expansion, mode index n corresponds to a y -direction wavenumber

$$k_y^{(n)} = \frac{2\pi n}{L_y}. \quad (\text{G.36})$$

The discrete modal separation q_0 therefore corresponds to a wavenumber difference

$$\Delta k_y = k_y^{(n+q_0)} - k_y^{(n)} = \frac{2\pi(n+q_0)}{L_y} - \frac{2\pi n}{L_y} = \frac{2\pi q_0}{L_y}. \quad (\text{G.37})$$

Using the wave relation $k = 2\pi/\lambda$, the spatial wavelength associated with this mode coupling is

$$\lambda_{\text{coupling}} = \frac{2\pi}{\Delta k_y} = \frac{2\pi}{2\pi q_0/L_y} = \frac{L_y}{q_0} = T. \quad (\text{G.38})$$

When the stripe period T matches this coupling wavelength, we achieve resonant enhancement analogous to the Bragg condition for coherent scattering in periodic structures (Kushwaha et al., 1993): reflections from successive interfaces interfere constructively, maximizing energy transfer between the mode pair (m, n) and $(m, n+q_0)$.

The period $T = L_y/q_0$ is not merely sufficient but necessary for fundamental-harmonic coupling. Any other period would either shift the fundamental component to a different mode index or distribute energy across multiple harmonics, reducing coupling strength at q_0 .

Part III: Amplitude maximization through square-wave profile. Given the resonant period $T = L_y/q_0$ from Part II, we now optimize the profile shape $f(y)$ to maximize $|\hat{V}_{q_0}|$ under constraint $|f(y)| \leq A$.

For a function $f(y)$ with period T , we can compute \hat{V}_{q_0} by integrating over one period and accounting for the number of periods in $[0, L_y]$. Since $f(y)$ repeats q_0 times over $[0, L_y]$, we have

$$\hat{V}_{q_0} = \frac{1}{L_y} \int_0^{L_y} f(y) e^{i2\pi q_0 y/L_y} dy = \frac{q_0}{L_y} \int_0^T f(y) e^{i2\pi q_0 y/L_y} dy. \quad (\text{G.39})$$

Substituting $T = L_y/q_0$ and changing variables $\eta = q_0 y/L_y$ (so $dy = T d\eta$ and $\eta \in [0, 1]$ over one period):

$$\hat{V}_{q_0} = \frac{q_0}{L_y} \int_0^1 f(T\eta) e^{i2\pi\eta} T d\eta = \frac{1}{L_y/q_0} \int_0^{L_y/q_0} f(y) e^{i2\pi q_0 y/L_y} dy. \quad (\text{G.40})$$

Now we apply Lemma G.3. Define a rescaled Fourier coefficient for one period:

$$\tilde{V} := \frac{1}{T} \int_0^T f(y) e^{i2\pi q_0 y / L_y} dy = \frac{1}{T} \int_0^T f(y) e^{i2\pi y / T} dy, \quad (\text{G.41})$$

where we used $q_0 / L_y = 1/T$. This is precisely the Fourier coefficient at the fundamental frequency for a function with period T , matching the form in Lemma G.3 with $L = T$ and $q = 1$.

From Lemma G.3, $|\tilde{V}|$ is maximized when $f(y)$ is the square wave

$$f(y) = A \cdot \text{sgn} \left[\cos \left(\frac{2\pi y}{T} \right) \right] = A \cdot \text{sgn} \left[\cos \left(\frac{2\pi q_0 y}{L_y} \right) \right], \quad (\text{G.42})$$

achieving

$$|\tilde{V}|_{\max} = \frac{2A}{\pi \cdot 1} = \frac{2A}{\pi}. \quad (\text{G.43})$$

From Eq. (G.40), we have $\hat{V}_{q_0} = \tilde{V}$, therefore

$$|\hat{V}_{q_0}|_{\max} = \frac{2A}{\pi}. \quad (\text{G.44})$$

The optimality of the square-wave profile admits dual interpretations. Mathematically, the Cauchy-Schwarz inequality (used in proving Lemma G.3) shows that phase alignment with the Fourier kernel $e^{i2\pi q_0 y / L_y}$ while maintaining maximal amplitude $|f(y)| = A$ maximizes the overlap integral in Eq. (G.32).

Physically, the square wave creates binary impedance contrasts at interfaces $y = (2n + 1)T/2$ for integer n . At each interface, waves reflect with coefficient approaching $R \rightarrow \pm 1$ as shown in Lemma G.1. When interfaces are spaced at period $T = L_y / q_0$, reflections from all interfaces arrive in phase at the coupled mode pair, creating maximal constructive interference. The Fourier coefficient $|\hat{V}_{q_0}|$ quantifies precisely this coherent amplification.

Synthesis: Joint optimization via amplitude and phase. Combining Parts I–III, the striped square wave

$$\delta c^2(x, y) = A \cdot \text{sgn} \left[\cos \left(\frac{2\pi q_0 y}{L_y} \right) \right] \quad (\text{G.45})$$

simultaneously achieves: (i) maximal modal coupling amplitude $|C_{(m,n),(m,n+q_0)}| = k_{m,n+q_0}^2 \cdot 2A/\pi$, and (ii) maximal spatial differentiation with dynamic range $\Delta c = 2A$. The latter creates maximal spatial variation in local resonance frequencies $\omega_{\text{res}}(y) \propto c(y)$ and consequently in frequency-dependent phase delays $\tau_i(y)$ (Theorem 2.1). Different input frequencies thus produce spatially distinct phase patterns, enabling frequency-selective detection. This joint optimization of coupling amplitude and phase selectivity maximizes drift detection sensitivity. \square

H. Prompt

The prompt used to make the classification is shown in Fig. H.1. To ensure a rigorous evaluation on the XD-Violence dataset (Wu et al., 2020), we adopt the ground-truth label definitions provided by the original benchmark authors.

A significant challenge in applying zero-shot Audio Language Models (ALMs) to this dataset is the nature of the source material; the majority of audio scenarios in XD-Violence are extracted from action movies and films. We observed that without domain-specific pre-training or fine-tuning, general-purpose ALMs tend to exhibit a “fictional bias.” Specifically, the models often recognize the high-energy acoustic signatures or specific dialogue as being part of a movie or video game and, as a result, classify the sample as containing “no violence” because it is perceived as non-real.

To mitigate this, we designed a specific Fictional Context Rule within our prompt. This rule explicitly instructs the model to ignore any mentions of the audio being from a fictional medium and to focus solely on the physical actions described (e.g., classifying a movie gunshot as a valid shooting event). This ensures that the ALM functions as a security-oriented classifier rather than a content-type identifier, allowing for a fair comparison against supervised baselines.

XD-Violence Audio Track Classification Prompt

[ROLE]

You are a professional security audio classifier.
Your ONLY task is to detect and categorize violent events from audio descriptions.

[EVIDENCE]

{evidence_text}

[RULES]

1. Label mapping (mandatory):

- B1:** Fighting (scuffles, hitting, physical combat)
- B2:** Shooting (gunfire, gunshots)
- B4:** Riot (large crowd chaos, protests, shouting)
- B5:** Abuse (harassment, physical or verbal abuse)
- B6:** Car accident (crashes, breaking glass, vehicle impacts)
- G:** Explosion (blasts, loud booming sounds)

None: Only if absolutely no violent events are mentioned.

2. Violence dominance rule:

If any single part of the evidence mentions a violent event, you MUST classify the sample based on that event, even if all other parts of the evidence state “there is no violence.” One positive detection overrides all **None** detections.

3. Fictional context rule:

Strictly IGNORE any mentions of the audio being from a “video game,” “movie,” “film,” or “soundtrack.” You must classify based on the actions described (e.g., if a movie gunshot is described, label it **B2**).

[OUTPUT FORMAT]

Output ONLY the code (e.g., **B4**). Do not include reasoning, descriptions, punctuation, or extra text.

Figure H.1. Prompt used for violence event classification from audio descriptions.

I. Additional Qualitative Comparisons

To complement the illustrative cases in Section 4.2, we provide six additional Mel-spectrogram examples of OWM. These examples highlight diverse acoustic scenarios and further demonstrate the relative robustness of OWM.

Car horn onset (Example R0003). As shown in Fig. I.1, a car horn appears at 21 s and persists for several seconds. Both methods successfully detect the onset, but OWM localizes the event with fewer spurious triggers.




Figure I.1. Example R0003: Car horn onset. OWM localizes events with few spurious triggers.

Intermittent piano playing (Example R0007). As shown in Fig. I.2, piano notes occur intermittently, producing pauses between onsets. OWM registers four discrete activations, thereby reducing sensitivity to the continuous drift caused by spectral divergence.




Figure I.2. Example R0007 (Chalet du Mont-Royal, Montreal — street traffic): The most salient event is intermittent piano playing, which triggers OWM four times; this discretization reduces sensitivity to the continuous drift caused by spectral divergence.

Conversational speech (Example R0028). As shown in Fig. I.3, after 30 s, a dialogue with pauses and speaker changes begins. OWM detects a single event at 38 s, reflecting its robustness to intra-speech variability.




Figure I.3. Example R0028 (Heping Road, Tianjin — street traffic): The most salient event is conversation after 30 s, with pauses and speaker changes. OWM is triggered only once at 38 s.

Repeated car horn sounds (Example R0030). As shown in Fig. I.4, horns occur around 24 s, 32 s, and 50 s. OWM detects only the late instance (around 51 s).




Figure I.4. Example R0030 (Century Clock, Tianjin — street traffic): The most salient event is repeated car horn sounds at approximately 24 s, 32 s, and 50 s. OWM detects only the instance around 51 s.

Railway station announcement (Example R0031). As shown in Fig. I.5, the segment contains continuous speech without other salient events. OWM triggers only once during the initial threshold adaptation phase, demonstrating its robustness against continuous background speech.




Figure 1.5. Example R0031 (Tianjin Railway Station, Tianjin — public square): Analysis of a railway station announcement. OWM triggers only once during the initial threshold adaptation, demonstrating its stability and resistance to false activations during continuous speech.

Church bell with crowd talking (Example R0131). As shown in Fig. 1.6, the spectrogram is dominated by continuous bell ringing. OWM’s adaptive threshold rises in response to this persistent signal, effectively suppressing background crowd noise and preventing spurious detections.




Figure 1.6. Example R0131 (Town Hall Square, Vilnius — square): OWM response to continuous church bell ringing. The adaptive threshold scales according to the persistent bell signal, effectively suppressing background crowd noise and preventing spurious detections.

Taken together, these additional examples reinforce the primary findings: OWM maintains high specificity by avoiding spurious detections during transient pauses, repeated motifs, or continuous sources, while remaining sensitive to salient novel events.

J. Additional FFT Analysis




Figure J.1. Frequency distribution across the p -field neurons obtained via FFT for different Examples (R0002, R0010, R0016, R0037, R0056). The results reveal band-specific activity in the theta (4–8 Hz), alpha (8–12 Hz), beta (13–30 Hz), and gamma (30–50 Hz) ranges.

In the main paper (Section 4.3), we presented FFT analyses for Examples R0016 and R0056 to illustrate how BioWM reallocates oscillatory activity around drift onsets. To provide further evidence, Fig. J.1 shows additional FFT spectra for five

1595 representative cases (R0002, R0010, R0016, R0037, R0056). These examples span a variety of acoustic scenes, including
1596 novel sound events (R0002), subcategory-level drift (R0010), transient pauses (R0016, R0037), and salient novel sources
1597 (R0056).

1598 Across all examples, the dominant oscillatory activity of p -field neurons lies within the θ (4–8 Hz), α (8–12 Hz), β (13–30
1599 Hz), and low- γ (30–50 Hz) bands, consistent with canonical neural regimes. Rather than being uniformly distributed,
1600 activity is organized in clustered regions that evolve after drift onsets. The figures display the activity of all p neurons over
1601 entire segments; thus, direct saliency patterns are not immediately visible. What becomes evident instead is the periodic
1602 coupling structure along the Y-axis of the lattice, in line with the striped optimality predicted by Theorem 2.4. This suggests
1603 that BioWM’s drift sensitivity arises from frequency-specific clustering combined with spatial coupling, rather than from
1604 broad or diffuse spectral fluctuations. Animated visualizations of these dynamics are included as supplementary material.
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649