Title: Semantic Flux: A Physics-Anchored Measure of Persistent Change for Pre-Linguistic AI Systems

by Stephen Francis, Phocoustic, Inc. January 30, 2026

Version 1.0, This document is intended as a technical white paper and may be updated in future versions.

Abstract

Conventional vision and artificial intelligence systems often conflate transient variation with meaningful change, leading to sensitivity without repeatability or reliable early warning capability. In physics-anchored perception architectures, measurement of change precedes interpretation and serves as a pre-decisional layer that determines whether downstream reasoning is warranted. This work formalizes **semantic flux** as a measurable quantity representing the accumulation of spatially localized, temporally persistent change after geometric relevance weighting and lineage enforcement. To characterize where this accumulation becomes concentrated, **semantic activation density** is introduced as a normalized indicator that distinguishes emergent structure from diffuse variation.

Semantic flux and activation density are treated as pre-decisional artifacts: they encode admissible change without assigning labels, classifications, or inferred meaning. Recurring drift structures give rise to **symbolic carriers of structured change**, compact representations that preserve how change evolves over time while remaining independent of language. Language models and other inference-based systems are positioned strictly downstream, consuming validated symbolic carriers to produce human-readable descriptions without participating in detection, filtering, or validation.

Evaluations on representative industrial inspection sequences illustrate improved repeatability, enhanced rejection of nuisance variation, and earlier detection of emergent anomalies compared to conventional change-based and inference-driven methods. By formalizing a measurable layer between raw perception and interpretation, this work provides an auditable and reliable foundation for downstream semantic reasoning grounded in persistent physical change.

1. Introduction

This paper describes a general measurement framework. It does not disclose implementation details, algorithms, or system architectures.

1.1 Motivation

Human perception exhibits a consistent and well-documented asymmetry: observers often sense that *something is changing* before they are able to determine *what that change represents*. This pre-interpretive sensitivity enables early awareness of instability, anomaly, or emergence even when visual evidence is weak, incomplete, or ambiguous. In many practical settings—industrial inspection, degraded sensing environments, or early-stage fault formation—this capability is critical.

By contrast, most contemporary artificial intelligence and computer vision systems bypass this stage entirely. They operate primarily on static representations or instantaneous differences, applying inference or classification directly to raw data. As a result, such systems tend to be highly sensitive to variation while remaining unreliable in the presence of noise, nuisance effects, or subtle, slowly developing change. The consequence is a familiar failure mode: systems that react strongly to transient disturbances yet fail to provide stable, repeatable early warning of meaningful change.

This gap motivates the need for a distinct measurement layer that precedes interpretation—one that determines whether change is *admissible* for semantic consideration before labels, explanations, or decisions are applied.

1.2 Key Observation

A central observation underlying this work is that **static scenes are informationally inert**. When a system observes no persistent change, there is no basis for interpretation beyond confirming stability. Conversely, not all change is meaningful. Transient variation, stochastic noise, and global nuisance effects may produce large instantaneous differences without conveying reliable information about system state or emerging structure.

Meaningful change arises only when variation is **persistent, localized, and structured across time**. Such change exhibits continuity, coherence, and lineage: it survives temporal filtering, concentrates spatially, and evolves in a manner that distinguishes it from random fluctuation. This observation suggests that meaning does not originate from instantaneous measurements or from inference alone, but from the accumulation of admissible change over time.

This work adopts the position that identifying and measuring this class of change is a prerequisite for reliable interpretation.

1.3 Thesis Statement

This paper proposes **semantic flux** as a measurable, repeatable quantity that captures admissible, structured change prior to symbolic or linguistic interpretation. Semantic flux is defined as the accumulation of spatially localized, temporally persistent change under explicit geometric and lineage constraints. It exists as a pre-decisional artifact: a measurement-layer construct that determines when downstream interpretation is warranted, without assigning labels or inferred meaning.

Importantly, semantic flux is presented as a **general measurement framework**, not as a system-specific feature. It is intended to be applicable across sensing modalities and architectural implementations. Phocoustic is introduced later in this work as one concrete instantiation that demonstrates the practical utility of semantic flux within a physics-anchored perception system, but it does not bound or define the framework itself.

1.4 Contributions

The primary contributions of this work are as follows:

- **Formal definition of semantic flux** as an accumulative, pre-interpretive measure of persistent change distinct from instantaneous variation or inference-based scores.
- **Introduction of semantic activation density**, enabling spatial localization and concentration analysis of admissible change.
- **A geometry-weighted, tiled measurement framework** that enforces locality and supports repeatable evaluation across time.
- **Explicit separation of detection and interpretation**, positioning semantic flux as a measurement-layer construct and downstream language or inference systems as consumers rather than generators of meaning.
- **Empirical validation against baseline methods**, demonstrating improved repeatability, noise rejection, and earlier detection of emergent anomalies in inspection scenarios.

Together, these contributions establish semantic flux as a transferable foundation for systems that require reliable early warning, explainability, and auditability, independent of any specific product, sensing modality, or interpretive mechanism.

2. Related Work

The problem of detecting meaningful change in visual data has been approached from multiple directions, including pixel-level differencing, statistical learning, physics-informed modeling, and, more recently, language-augmented vision systems. Each contributes useful tools, but none directly address the core question posed in this paper: how to define and measure *change that persists* in a principled, modality-agnostic way.

2.1 Frame Differencing and Optical Flow

Classical techniques such as frame differencing, background subtraction, and optical flow are explicitly designed to detect change between successive frames. These methods are computationally efficient and sensitive to small variations in motion or intensity. However, they are inherently local and instantaneous. Noise, illumination fluctuation, and viewpoint jitter are often indistinguishable from meaningful change.

Crucially, these approaches lack persistence and lineage. Change is detected, but not remembered. There is no mechanism to determine whether a deviation represents a transient fluctuation, a stable transformation, or the early stages of a developing anomaly. As a result, sensitivity is achieved at the cost of reliability.

2.2 Statistical and CNN-Based Anomaly Detection

Statistical models and convolutional neural networks (CNNs) dominate modern anomaly detection in vision. These systems excel at pattern recognition when trained on large, representative datasets. When anomalies closely resemble those seen during training, performance can be strong.

However, such systems are fundamentally reference-dependent. They require prior examples, retraining for new domains, and careful curation to avoid bias. Generalization outside the training distribution remains fragile, particularly for low-contrast, emergent, or previously unseen changes. Failure modes are often opaque: a misclassification provides little insight into *why* a change was ignored or amplified.

2.3 Physics-Informed Vision

Physics-informed approaches attempt to ground perception in constraints such as geometry, optics, material behavior, or energy conservation. These methods improve interpretability and reduce some forms of spurious detection. They are especially effective in controlled environments where physical models are well understood.

Nonetheless, most physics-informed systems remain state-based. They describe what *is* observed at a given moment, rather than how admissible change accumulates over time. Temporal persistence is often treated as a secondary filter rather than as a first-class quantity.

2.4 Language Models in Vision Systems

Recent vision-language systems leverage large language models to explain, summarize, or contextualize visual inputs. These models are powerful at interpretation and communication, particularly when reasoning about objects, scenes, or actions.

Yet language models are weak detectors. They rely on upstream perception systems to provide stable inputs and are often conflated with the task of measurement itself. Without a grounded representation of persistent change, linguistic explanations risk being confident descriptions of unstable or noisy signals.

2.5 Summary of Gaps

Across these approaches, a common limitation emerges. There is no explicit, general measure of change that survives time—one that distinguishes transient variation from structurally meaningful evolution. Moreover, there is no principled boundary separating *measurement* from *interpretation*. Detection and explanation are frequently entangled, obscuring failure modes and limiting generality.

The Semantic Flux framework is proposed to address these gaps by treating persistent change as a measurable quantity in its own right, prior to and independent of symbolic or linguistic interpretation.

3. Conceptual Framework

The Semantic Flux framework reframes perception around a single organizing principle: *persistent change*, rather than instantaneous state. This section introduces the conceptual assumptions that govern the measurement of semantic flux and motivate its separation from interpretation.

3.1 Stability as the Null Hypothesis

In the proposed framework, stability is treated as the null hypothesis. A scene that remains invariant over time—within physically admissible tolerances—carries no semantic signal. Static structure, regardless of visual complexity, is informationally inert with respect to change.

This stands in contrast to many vision systems that attempt to extract meaning from single frames or fixed configurations. Semantic flux instead assumes that meaning does not reside in structure alone, but in the *departure* from structure. If nothing is changing, there is nothing to explain, predict, or interpret.

Under this assumption, the absence of persistent change is not a failure case; it is a valid and expected outcome corresponding to semantic zero.

3.2 Change vs. Noise

Not all variation constitutes meaningful change. The framework distinguishes sharply between noise and drift.

Noise is characterized by being transient, incoherent, and non-local. It appears sporadically, lacks directional consistency, and does not maintain identity across time. Examples include sensor jitter, illumination flicker, compression artifacts, or stochastic pixel fluctuations. Noise does not accumulate and cannot support lineage.

Drift, by contrast, is persistent, localized, and directional. It exhibits continuity across frames, maintains spatial coherence, and evolves in a manner consistent with physical or structural constraints. Drift can be weak, gradual, or low-contrast, yet still meaningful if it survives temporal validation.

Semantic flux is defined only over drift. Noise is explicitly excluded not through heuristic thresholds, but through the requirement of persistence and directional consistency.

3.3 Field Interpretation

Rather than treating change as a sequence of isolated events, semantic flux models change as a discrete field defined over space and time. Each local region contributes to a flow of change, and these flows may strengthen, decay, or interact across frames.

In this view, meaning eligibility does not arise from any single state of the system. It emerges from the *flow*—the structured evolution of change across the field. A region becomes semantically interesting not because of what it looks like at an instant, but because of how it moves through admissible change space over time.

This field-based interpretation provides a natural boundary between measurement and interpretation. Semantic flux quantifies the flow. Language, labels, or decisions may act upon it later, but they are not required for its existence.

4. Formal Definitions

This section introduces the minimal formal machinery required to define semantic flux as a measurable quantity. The definitions are intentionally discrete, finite, and implementation-agnostic. No continuous field theory or probabilistic assumptions are required.

4.1 Spatial Discretization (Tiling)

Let an input scene be represented as a sequence of frames over time. Each frame is partitioned into a fixed grid of spatial tiles. These tiles serve as local measurement cells and define the atomic units over which change is evaluated.

Tiling imposes no semantic meaning by itself. Its purpose is to localize measurement, constrain spatial support, and enable consistent temporal comparison. All subsequent operators act on tile-indexed quantities.

4.2 Geometry-Weighted Drift

Not all spatial regions contribute equally to a given measurement. The region of interest is the subset of the observed scene selected according to geometric, optical, or structural constraints that determine where meaningful change is expected to occur.

Frame-to-frame change is first computed locally per tile, then projected onto the region of interest. Tiles that fall outside the ROI, or that violate geometric admissibility constraints, are suppressed or down-weighted.

This step enforces locality and prevents diffuse, scene-wide variation from contributing to semantic flux. Only change that is geometrically consistent with the region under evaluation is retained.

4.3 Persistence Operator

Let WWW denote a finite temporal window spanning multiple frames. A persistence operator is applied to the tile-wise change signals across this window.

For change to be considered admissible, it must:

- Persist across the window WWW,
- Maintain directional coherence (i.e., not oscillate randomly), and
- Remain localized to a consistent spatial neighborhood.

Transient, incoherent, or reversing signals are rejected at this stage. Persistence is therefore not a smoothing operation, but a validation gate that separates drift from noise.

4.4 Semantic Flux

Semantic flux is defined as the accumulated admissible change within a region R over a time window T.

Semantic flux is computed by iterating over each frame in the time window and, at each frame, summing the tile-wise changes within the region of interest that have passed geometric filtering and temporal persistence validation. The final value reflects the total accumulated admissible change across both space and time.

No exotic mathematics are implied. Semantic flux is an additive quantity that increases only when admissible change accumulates over time.

4.5 Semantic Activation Density

To enable comparison across regions and time scales, semantic flux may be normalized by spatial extent and temporal duration.

Semantic activation density expresses how concentrated persistent change is, on average, per unit area and per unit time. It is obtained by dividing the accumulated admissible change by the spatial extent of the region being evaluated and by the duration of the observation window.

where $|R|$ denotes the number of spatial tiles in the region of interest and $|T|$ denotes the number of discrete time indices in the analysis window.

This quantity reflects the intensity of persistent change per unit area and per unit time. It allows small but highly active regions to be distinguished from large regions with diffuse or marginal drift.

Semantic activation density reflects the intensity of persistent change per unit area and per unit time, not a probabilistic or physical density.

4.6 Symbolic Carriers

Semantic flux does not directly encode meaning. Instead, recurring patterns of admissible drift may be assigned **symbolic carriers**—compact identifiers that reference characteristic modes of change.

Symbolic carriers do not represent appearance, texture, or object identity. They encode *how change evolves*: its directionality, persistence profile, and geometric footprint. These carriers enable downstream systems to reference, compare, and reason about recurring change structures without reprocessing raw measurements.

Crucially, symbolic carriers are optional and downstream. Semantic flux exists independently of any symbolic assignment.

5. System Architecture

The Semantic Flux framework is realized as a layered architecture that enforces a strict separation between measurement, symbol formation, and interpretation. Each layer operates on well-defined inputs and outputs, and no layer is permitted to subsume the role of another.

5.1 Measurement Layer (Pre-Linguistic)

The measurement layer operates entirely below language and symbolism. Its function is to detect, validate, and accumulate admissible change without assigning meaning or labels.

This layer includes:

- **Tiling**, which discretizes the scene into local measurement cells;
- **Geometry weighting**, which restricts contribution to geometrically admissible regions of interest;
- **Persistence filtering**, which enforces temporal survival and directional coherence; and
- **Flux accumulation**, which produces semantic flux values and derived densities.

All processing at this stage is deterministic, physically grounded, and time-aware. The result is a set of numerical measures that describe how change accumulates over specific regions and time intervals. No symbolic interpretation or semantic labeling is performed at this stage.

5.2 Symbol Formation Layer

The symbol formation layer converts validated measurements into compact, reusable representations without introducing language.

Its responsibilities include:

- **Drift clustering**, where recurring patterns of admissible change are grouped based on similarity of evolution;
- **Lineage tracking**, which maintains identity of drift structures across time windows; and
- **Carrier assignment**, where stable drift patterns are assigned symbolic carriers.

These carriers function as references to *how* change behaves, not *what* it represents. They encode persistence profiles, geometric footprint, and directional evolution. The result is a symbolic substrate that is stable, auditable, and detached from raw sensory data.

5.3 Interpretation Layer (LLM)

The interpretation layer consumes symbolic carriers and associated metadata to produce human-readable explanations, summaries, or decisions.

Its inputs consist exclusively of:

- Symbolic carrier identifiers,
- Temporal and spatial metadata, and
- Quantitative flux descriptors.

Its output is language.

Critically, this layer has **no access to raw sensor data, tiles, frames, or flux computation mechanisms**. It cannot influence measurement, persistence validation, or symbol formation.

Key Claim

Language models do not and cannot generate semantic flux.

Semantic flux arises only from persistent, admissible change measured over space and time. Language models operate solely on symbolic inputs provided to them. They may explain, contextualize, or reason about flux-derived symbols, but they cannot create, amplify, or suppress semantic flux itself.

This architectural separation ensures that detection remains grounded, interpretation remains accountable, and failure modes are observable rather than entangled.

6. Experimental Design

The experimental design is constructed to evaluate whether semantic flux can reliably distinguish persistent, meaningful change from transient variation across representative inspection scenarios. The emphasis is on longitudinal behavior rather than single-frame accuracy.

6.1 Datasets

Three classes of image sequences are used to evaluate the framework:

- **PCB inspection sequences**, consisting of repeated captures of printed circuit board regions under nominally stable conditions. These sequences include fine traces, solder features, and low-contrast surfaces where early defects are difficult to detect visually.
- **Wafer surface sequences**, comprising tiled or cropped views of semiconductor wafer regions captured across time. These sequences emphasize subtle surface non-uniformities, faint line structures, and slow material evolution.
- **Controlled micro-change overlays**, where synthetic or semi-synthetic perturbations are introduced into otherwise stable sequences. These overlays simulate faint, localized changes that evolve gradually across frames while remaining below visual salience thresholds.

Together, these datasets span real-world inspection data and controlled test cases, allowing both qualitative and quantitative evaluation.

Detailed datasets are omitted here and will be included in future technical or journal versions of this work

6.2 Conditions

Each sequence is evaluated under one of three controlled conditions:

1. **Stable**
   No physically meaningful change is present. Minor sensor noise, illumination variation, or compression artifacts may occur, but no persistent drift is introduced.
2. **Nuisance-only**
   Sequences include transient disturbances such as flicker, jitter, or global intensity fluctuation. These effects are designed to challenge sensitivity while lacking persistence or spatial coherence.
3. **Emergent micro-defect**
   A small, localized change evolves gradually across time. The change is persistent, directional, and spatially consistent, but may remain visually indistinguishable in individual frames.

These conditions are designed to test the null hypothesis of stability, the rejection of noise, and the detection of admissible drift, respectively.

6.3 Baselines

Semantic flux measurements are compared against common change-detection and anomaly-detection baselines, including:

- **Frame difference energy**, computed as the summed absolute pixel differences between successive frames;
- **Structural Similarity Index (SSIM) delta**, measuring perceptual structural change between frames;
- **Optical flow magnitude**, aggregated over the region of interest; and
- **CNN-based anomaly score** (optional), derived from a trained convolutional model where applicable.

Baselines are evaluated using identical regions and time windows to ensure comparability. Performance is assessed in terms of false activation under stable and nuisance-only conditions, and early activation under emergent micro-defect conditions.

7. Metrics

Evaluation focuses on temporal reliability, noise discrimination, and spatial consistency rather than single-frame accuracy. All metrics are computed over repeated runs using identical data and windowing parameters unless otherwise noted.

Detailed numerical tables are omitted here and will be included in future technical or journal versions of this work.

7.1 Repeatability (Coefficient of Variation)

Repeatability measures the stability of semantic flux outputs across identical experimental runs. For a fixed sequence and region of interest, semantic flux values are computed multiple times under the same conditions.

Repeatability is quantified using the **coefficient of variation (CV)**, defined as the ratio of the standard deviation to the mean of the measured flux values. Lower CV indicates higher repeatability and reduced sensitivity to stochastic variation.

This metric evaluates whether semantic flux behaves as a stable measurement rather than a volatile score.

7.2 Noise Rejection Ratio

The noise rejection ratio compares semantic activation under nuisance-only conditions to activation under stable conditions.

For each sequence class, the ratio is computed as the mean semantic flux (or activation density) observed under nuisance perturbations divided by that observed under stable sequences. Ratios near unity indicate effective suppression of nuisance variation.

This metric assesses the framework's ability to treat noise as non-semantic without requiring explicit noise modeling.

7.3 Persistence Lift

Persistence lift measures the degree to which detections survive temporal validation.

It is defined as the fraction of detections that remain active across a predefined persistence window relative to the total number of initial activations. Higher values indicate that detected changes are temporally coherent rather than transient spikes.

Persistence lift directly reflects the effectiveness of the persistence operator in distinguishing drift from noise.

7.4 ROI Localization Consistency

ROI localization consistency evaluates spatial stability of detected change.

For each sequence, the top-k regions of highest semantic flux are identified per time window. Consistency is measured as the fraction of these regions that remain within the same spatial neighborhood across successive windows.

This metric quantifies whether detected change maintains spatial identity over time, rather than wandering due to noise or global effects.

7.5 Lead Time

Lead time measures the temporal advantage of semantic flux detection relative to baseline methods.

For emergent micro-defect sequences, lead time is defined as the number of frames by which semantic flux activation precedes the first reliable detection by baseline metrics (e.g., SSIM delta, optical flow magnitude, or CNN anomaly score).

Positive lead time indicates earlier recognition of persistent change, even when the change is not yet visually salient.

8. Results

Results are presented to illustrate the temporal, spatial, and repeatability characteristics of semantic flux under the experimental conditions described in Section 6. No task-specific tuning or post hoc thresholding was applied beyond parameters fixed prior to evaluation.

8.1 Time-Series Comparisons

Time-series plots of semantic activation were generated for all sequences, comparing geometry-weighted semantic flux against unweighted change accumulation.

Across stable and nuisance-only conditions, unweighted measures exhibited frequent low-level activation driven by global variation and transient noise. Geometry-weighted flux remained near baseline, showing minimal accumulation over time.

Under emergent micro-defect conditions, geometry-weighted semantic flux exhibited a gradual, monotonic increase consistent with persistent localized change. Unweighted measures either responded late or showed oscillatory behavior that did not accumulate reliably.

These comparisons demonstrate that weighting by region geometry materially alters temporal behavior, suppressing diffuse variation while preserving persistent drift.

8.2 Spatial Localization

Spatial distributions of semantic flux were visualized as tile-wise heatmaps overlaid on the original frames.

In stable and nuisance-only sequences, activation was sparse and spatially inconsistent, with no tile retaining elevated flux across windows. In emergent micro-defect sequences, activation localized to a compact region and intensified gradually over time.

Importantly, the location of peak activation remained stable across successive windows, even when the underlying visual change remained difficult to perceive in individual frames. This stability was not observed in baseline spatial measures.

8.3 Repeatability Tables

Repeatability metrics were summarized in tabular form across multiple identical runs.

Semantic flux measurements exhibited low coefficients of variation across all sequence classes, with the lowest variability observed in stable and nuisance-only conditions. Emergent micro-defect sequences showed slightly higher variance, attributable to gradual signal accumulation, but remained well within acceptable bounds for longitudinal measurement.

Baseline methods showed substantially higher variability, particularly in nuisance-only conditions, where transient effects produced inconsistent activations across runs.

Detailed numerical tables are omitted here and will be included in future technical or journal versions of this work.

8.4 Early Warning Examples

Representative early warning cases are presented as visual sequences paired with activation plots.

In these examples, semantic flux crossed activation thresholds several frames prior to any reliable indication from baseline methods. Visual inspection confirmed that the underlying change was present but not salient at the time of activation.

These cases illustrate that semantic flux responds to the *persistence* of change rather than its immediate visibility, providing early indication without reliance on trained templates or appearance models.

9. Discussion

The results presented in Section 8 highlight a consistent pattern: semantic flux behaves as a stable measurement of persistent change, while baseline methods tend to respond to instantaneous variation. This section explains why the framework succeeds, where its limits lie, and why language-only systems cannot substitute for it.

9.1 Why Semantic Flux Works

Semantic flux succeeds because it enforces three constraints that are typically violated or weakened in conventional vision systems.

First, it enforces **locality**. Change is evaluated within bounded spatial tiles and projected onto defined regions of interest. This prevents diffuse, scene-wide variation from accumulating semantic weight and ensures that activation remains tied to specific spatial structures.

Second, it enforces **time**. Change must survive across a persistence window and maintain directional coherence. Instantaneous differences, regardless of magnitude, do not qualify. This requirement converts detection from a snapshot problem into a longitudinal measurement.

Third, it enforces **geometry**. Only change that is consistent with the geometry of the region under evaluation contributes to semantic flux. This constraint filters out changes that are spatially inconsistent with the underlying structure, even if they are visually prominent.

Together, these constraints ensure that semantic flux accumulates only when change is physically plausible, spatially coherent, and temporally persistent.

9.2 Failure Modes

Semantic flux is not intended to detect all forms of change.

One failure mode arises under **global illumination collapse**, such as abrupt lighting loss or saturation affecting the entire scene uniformly. In such cases, locality and geometry weighting may suppress activation, correctly interpreting the event as non-structural.

Another limitation occurs with **extremely rapid catastrophic change**, where meaningful transformation happens within fewer frames than the persistence window allows. In these cases, semantic flux may lag detection by design, favoring reliability over immediacy.

These failure modes reflect deliberate design trade-offs rather than implementation defects.

9.3 Why LLMs Alone Cannot Do This

Large language models are powerful tools for explanation and reasoning, but they cannot replace semantic flux.

Language models do not maintain **persistence memory** over raw sensory data. They operate on provided tokens, not on evolving physical signals across time.

They lack **physical grounding**. Without access to geometry, spatial locality, and admissible change constraints, they cannot distinguish noise from drift in a principled way.

They also lack **locality constraints**. Language models process symbols globally; they do not enforce spatial neighborhood consistency or region-specific validation.

As a result, language models may describe, summarize, or contextualize semantic flux, but they cannot generate it. Semantic flux must exist prior to language.

10. Limitations

The Semantic Flux framework is intentionally constrained. Its strengths arise from these constraints, but they also define clear limitations.

First, semantic flux **requires temporal data**. Because it measures persistent change, it cannot operate on single images or isolated snapshots. Applications that lack repeated observation over time are outside its scope.

Second, the framework **requires region-of-interest geometry**. Locality and geometric weighting are core to noise rejection and drift validation. When no meaningful geometric constraints can be defined, semantic flux may become overly conservative or ambiguous.

Third, semantic flux is **not a replacement for classifiers**. It does not assign object identity, defect class, or semantic labels. Instead, it provides a pre-linguistic measurement of change that may be consumed by downstream classifiers or decision systems.

Finally, semantic flux does **not constitute semantic "understanding."** It does not reason, infer intent, or interpret meaning. It measures admissible change and nothing more. Any notion of understanding arises only when semantic flux is combined with symbolic or linguistic interpretation layers.

These limitations are deliberate. By restricting scope, the framework maintains reliability, auditability, and generality across domains.

11. Broader Implications & Generality

Semantic flux is presented in this work as a general measurement framework for persistent change, rather than as a system-specific feature or product-bound capability. The core contribution lies in formalizing an intermediate quantity that distinguishes admissible, structured change from transient variation prior to semantic labeling or inference. By separating measurement from interpretation, the framework addresses a foundational challenge shared across perception, inspection, and artificial intelligence systems: determining when downstream reasoning is warranted.

The framework is intentionally modality-agnostic. While experimental demonstrations in this work derive semantic flux from visual inspection sequences, the underlying principles of spatial localization, temporal persistence, and lineage consistency apply equally to other time-varying signals, including acoustic measurements, electromagnetic sensing, and hybrid or structured illumination modalities. Semantic flux operates on normalized representations of change rather than raw sensor values, enabling transfer across domains without requiring domain-specific retraining or semantic priors.

From a systems perspective, semantic flux occupies a pre-decisional measurement layer. It produces auditable artifacts that characterize how change accumulates and concentrates over space and time without assigning labels, classifications, or inferred meaning. This positioning allows semantic flux to complement, rather than replace, existing inference-based or learning-based methods. By constraining interpretation to regions and intervals where persistent change is present, the framework can reduce false positives, improve repeatability, and support earlier detection of emergent phenomena.

Semantic flux provides a missing measurement layer between raw perception and language. Rather than treating meaning as an emergent property of static scenes or model inference, it frames meaning eligibility as a function of persistent, localized change. By enforcing locality, geometry, and time, semantic flux transforms detection into a repeatable measurement problem: change is validated before it is interpreted, accumulated before it is labeled, and bounded before it is explained. This enables systems to respond predictively to emerging structure without relying on training data, appearance models, or linguistic inference at the detection stage.

Phocoustic serves as one implementation that demonstrates the practical utility of semantic flux within a physics-anchored perception architecture. In this context, semantic flux supports early anomaly detection and downstream interpretability while remaining independent of the specific mechanisms used for explanation or reporting. However, the framework itself is not tied to Phocoustic or to any particular software stack, sensing platform, or language model. Alternative implementations may employ different discretization strategies, persistence criteria, or downstream reasoning systems while preserving the core measurement principles described here.

More broadly, semantic flux contributes to ongoing efforts to improve reliability and transparency in intelligent systems by re-establishing measurement as a prerequisite for interpretation. Measurement produces evidence; language may act upon it—but cannot replace it. By formalizing persistent change as a measurable, transferable quantity, this work offers a foundation for safer, more reliable perceptual systems across industrial, environmental, and intelligent applications.

Appendices

Appendix A. Notational Conventions and Scope

This appendix clarifies the representational conventions used throughout the paper.

All quantities and operations described in the main text are discrete, finite, and expressed in plain language. Mathematical notation is used sparingly and only to support conceptual clarity. No continuous field assumptions, probabilistic models, or closed-form analytical solutions are required for the definitions presented.

Spatial references (such as regions of interest and tiles) are treated as bounded, finite partitions of an observed scene. Temporal references (such as time windows and persistence intervals) refer to finite sequences of discrete observations. Normalization and accumulation operations are described descriptively rather than symbolically to emphasize interpretation over formalism.

This choice is intentional. The objective of the paper is to define a measurement framework that is precise, auditable, and transferable across domains without requiring specialized mathematical machinery. Readers should be able to understand and apply the concepts of semantic flux and semantic activation density based on their definitions and constraints, independent of notation.

Appendix B. Tile Size Sensitivity Study

A sensitivity study was conducted to assess the impact of tile size on semantic flux behavior.

Smaller tiles increased spatial precision but introduced higher susceptibility to noise, requiring stronger persistence filtering. Larger tiles reduced noise sensitivity but degraded localization and diluted early micro-drift signals.

Across datasets, intermediate tile sizes produced the most stable trade-off between localization consistency, repeatability, and lead time. Importantly, semantic flux behavior remained qualitatively consistent across tile sizes, indicating robustness to discretization choice.

Tile size selection therefore represents a tunable resolution parameter rather than a failure point of the framework.

Appendix C. Conceptual Pipeline Description

This appendix provides a high-level conceptual description of the semantic flux measurement pipeline. It is intended to clarify processing stages rather than to specify algorithms or implementation details.

For each frame in a time sequence, the spatial domain is partitioned into local tiles, and a measure of local change is computed at each tile location relative to the preceding frame.

Change contributions are then evaluated with respect to region geometry. Only tiles that lie within the defined region of interest contribute to subsequent accumulation; change occurring outside the region is suppressed or ignored.

Change that passes geometric relevance is subjected to temporal persistence validation over a finite window. Only change that is directionally coherent and persists across time contributes to the admissible change signal.

Semantic flux is computed by accumulating admissible change across all tiles within a region and across all time indices within the analysis window. Semantic activation density is obtained by normalizing this accumulated change by the spatial extent of the region and the duration of the time window.

Symbol formation, lineage tracking, and interpretive processing occur downstream of semantic flux computation and are not part of the measurement process described here.

Appendix D. Additional Visualizations

Supplementary figures include:

- Extended time-series plots comparing semantic flux with baseline metrics
- Tile-wise heatmaps illustrating flux accumulation under nuisance and defect conditions
- Repeatability plots across identical experimental runs
- Early warning examples with synchronized frame sequences and activation curves

These visualizations reinforce the results presented in Section 8 and are provided to support qualitative inspection and reproducibility.