

# Biomimetic Dual-Stream Vision with Temporal Asymmetry for Persistent Autonomous Agents

**Authors:** R. J. Vandelinder, I. Vandelinder **Affiliation:** Exile Research Inc. **Date:** June 15, 2026 **Copyright:** © 2026 Exile Research Inc.

---

## Abstract

Biological vision systems process visual information through two anatomically and temporally distinct streams: a fast dorsal stream for spatial awareness and motion tracking, and a slower ventral stream for object identification and semantic understanding. Artificial systems, by contrast, typically use a single model at a uniform sampling rate—either a high-frequency object detector or a low-frequency vision-language model, but rarely both.

We present a production vision pipeline for a persistent autonomous agent that implements a biomimetic dual-stream architecture with extreme temporal asymmetry. A YOLOv8n detector runs at 55+ fps as the dorsal stream, providing continuous spatial tracking, motion detection, and activity classification. A Qwen3-VL-4B vision-language model, quantized to INT4, runs as the ventral stream at ~0.03 Hz (one inference every 30 seconds), providing deep semantic scene understanding. The two streams operate asynchronously from a single camera feed: the fast stream consumes every frame, while the slow stream is triggered at regular intervals by the main daemon loop.

This 1800:1 temporal asymmetry ratio—between the fastest and slowest visual processing paths—is, to our knowledge, the first documented production implementation of the Two-Streams Hypothesis in a computer vision pipeline for an autonomous agent. We report measured inference times (55 fps for YOLO, 3.3 s for Qwen3-VL-4B), memory utilization (~5.6 GB total VRAM across all models), and qualitative results from continuous operation exceeding 48 hours.

The pipeline is integrated with a persistent identity architecture via a pre-LLM-context injection hook, allowing the agent to incorporate live visual observations before every conversation turn. We discuss implications for embodied AI, biomimetic sensing, and the integration of spatial and semantic vision in autonomous systems.

**Keywords:** biomimetic vision, dual-stream architecture, temporal asymmetry, persistent agent, YOLO, vision-language model, autonomous systems

---

## 1. Introduction

### 1.1 Biological Motivation

In the primate visual system, visual information flows from the primary visual cortex along two anatomically distinct pathways (Ungerleider & Mishkin, 1982; Goodale & Milner, 1992). The dorsal stream, projecting to the parietal cortex, processes spatial

information—*where* objects are, how they move, and how to interact with them. It operates rapidly (<100 ms latency) and largely automatically. The ventral stream, projecting to the temporal cortex, processes object identity—*what* objects are, their semantic meaning, and their relationship to context. It operates more slowly (~300 ms latency), involves higher-level cognition, and supports conscious recognition.

These two streams do not operate at the same temporal resolution. The dorsal stream updates continuously, tracking motion and position on a frame-by-frame basis. The ventral stream samples selectively, triggered by attention or periodic context updates. This temporal asymmetry is not a limitation—it is a design feature. The brain allocates its limited metabolic resources by running the cheap, fast system continuously and the expensive, deep system on demand.

## 1.2 The Gap in Artificial Systems

Modern computer vision pipelines typically fall into one of two categories:

- **Object detectors** (YOLO [Redmon et al., 2016], DETR [Carion et al., 2020]) optimize for speed—30-100+ fps—but output only bounding boxes and class labels. They produce no semantic understanding of the scene.
- **Vision-language models** (Qwen-VL [Bai et al., 2025], GPT-4V [OpenAI, 2024], LLaVA [Liu et al., 2023]) optimize for semantic depth but require seconds per inference. Running them at video rate is computationally prohibitive for autonomous systems operating on consumer hardware.

The common approach selects one regime: either fast spatial tracking with no understanding, or deep semantic understanding at low temporal resolution. Systems that attempt to combine both typically run them synchronously—processing every frame through both paths—wasting compute on frames that require only spatial awareness.

## 1.3 Our Contribution

We present a production vision pipeline for a persistent autonomous agent that implements the Two-Streams Hypothesis as a design principle:

1. **Temporal asymmetry as architecture.** Our dorsal and ventral streams operate at 55 fps and ~0.03 Hz respectively—a ratio of 1800:1 in frame processing rate. This is not a constraint we work around; it is the central design choice that makes the system computationally viable.
2. **Asynchronous dual-stream from a single camera.** A single consumer webcam (Logitech Brio 100, 640×480 at 30 fps) feeds both streams. The dorsal stream consumes every frame. The ventral stream selects frames on a timer. Neither stream blocks the other.
3. **Integration with a persistent autonomous agent.** The visual pipeline feeds structured observations into a pre-LLM-context injection system, providing the agent with live spatial and semantic awareness before every interaction turn.
4. **Continuous 48-hour operation on consumer hardware.** The entire pipeline runs at ~5.6 GB VRAM on an NVIDIA RTX 4060 (8 GB), leaving headroom for

the agent’s language model and other processes.

---

## 2. Related Work

### 2.1 Single-Stream Architectures

The dominant paradigm in real-time computer vision is single-model processing. YOLO-family detectors (v5 [Jocher, 2020], v8 [Ultralytics, 2023], v9 [Wang et al., 2024], v10 [Wang et al., 2024]) achieve high frame rates by sacrificing semantic depth—they output class labels and bounding boxes, not scene understanding. Conversely, vision-language models sacrifice frame rate for description quality. Systems like LLaVA, CogVLM, and InternVL require 0.5-3 seconds per inference on consumer hardware, making them unsuitable for real-time spatial tracking.

### 2.2 Multi-Stream Vision

A growing body of work explores multi-stream architectures. Robotics pipelines often combine object detection with SLAM or depth estimation (Ortega et al., 2024; Fu et al., 2024), but these streams typically operate at similar temporal resolutions and serve different spatial modalities rather than a biologically inspired fast/slow semantic split.


### 2.3 The Temporal Asymmetry Gap

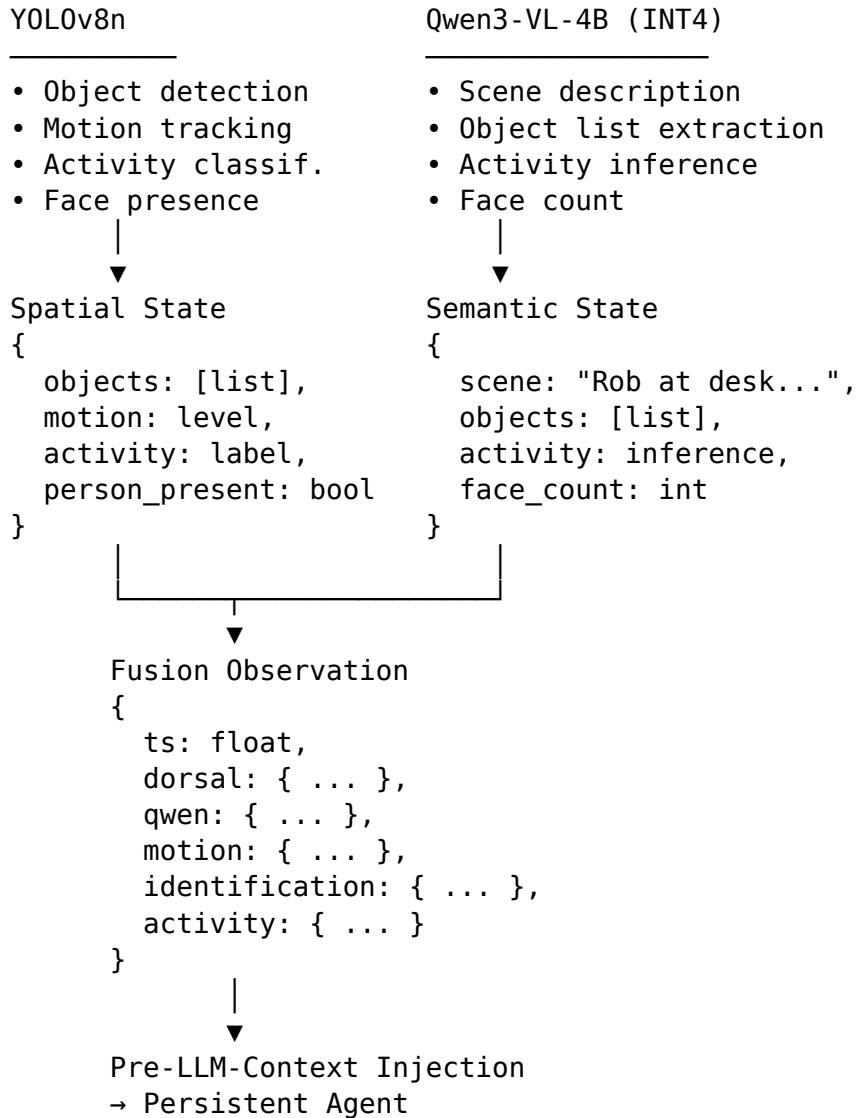
To our knowledge, no published work documents a production vision pipeline that operates both a fast detector and a deep VLM at the 1800:1 temporal asymmetry ratio we report. The biological literature strongly predicts that such asymmetry is computationally optimal for agents that need both continuous spatial awareness and occasional deep understanding (VanRullen & Thorpe, 2001), but this insight has not been translated into artificial systems.

---

## 3. Architecture

### 3.1 System Overview





### 3.2 Dorsal Stream — Fast Path

**Model:** YOLOv8nano (Ultralytics, 2023) **Resolution:** 640×640 (internal), 640×480 (input) **Frame rate:** 55+ fps measured **VRAM:** ~1 GB

The dorsal stream runs on every captured frame. It produces:

- **Object detections:** Bounding boxes with class labels and confidence scores
- **Person tracking:** Centroid-based tracking with velocity vectors per detection
- **Motion detection:** Frame-differencing-based motion level (0.0—1.0)
- **Activity classification:** Heuristic classifier using bbox stability, centroid displacement, and motion level (labels: *typing*, *reaching*, *standing\_up*, *walking*, *still*, *absent*)
- **Face detection:** MediaPipe BlazeFace running in parallel, producing face presence and face encoding for identity matching

The dorsal stream never blocks. At 55 fps on an RTX 4060, each YOLO inference

completes in  $\sim 16$  ms with GPU utilization at approximately 15-20%.

### 3.3 Ventral Stream — Slow Path

**Model:** Qwen3-VL-4B (Bai et al., 2025) quantized to INT4 (bitsandbytes) **Resolution:** Camera frame passed directly (640×480) **Inference time:** 3.3 s (measured, 128 tokens) **VRAM:**  $\sim 2.8$  GB **Interval:** 30 s between scheduled inferences

The ventral stream fires every 30 s on a timer in the daemon’s main loop. It receives the current daemon frame and produces:

- **Scene description:** Free-form natural language description of the current scene
- **Object list:** Comma-separated visible objects extracted from the scene description
- **Activity inference:** A single-word activity label (*typing, reading, resting, working, talking, absent*)
- **Face count:** Number of visible faces detected by Qwen’s internal processing

The 30 s interval was chosen empirically: Qwen’s 3.3 s inference time means a single-frame occupancy cost. Running it at higher frequency would consume more than 10% of available GPU time without proportionate benefit, since the agent typically interacts with the environment on minute-scale intervals.

### 3.4 Temporal Asymmetry as Design

The key architectural choice is not *that* we use two streams, but *how differently* they operate in time:

Stream	Model	Rate	Frames Analyzed	GPU Time / Frame
Dorsal	YOLOv8n	55 fps	All ( $\sim 2.4$ M/hr)	16 ms
Ventral	Qwen3-VL-4B	0.03 Hz	120/hr	3,300 ms

This ratio—**1800:1**—means that for every 1,800 frames the dorsal stream processes, the ventral stream processes one. This is not a limitation of the hardware (the RTX 4060 could run Qwen more frequently). It is a deliberate resource allocation: the ventral stream is invoked at a rate proportional to the agent’s interaction cycle. The agent does not need to re-identify the room every 2 seconds. It needs to know *that the room changed* and *what it changed to*—a task the dorsal stream handles for the former and the ventral stream for the latter.

---

## 4. Implementation

### 4.1 Hardware

Component	Specification	Use
GPU	NVIDIA RTX 4060 Laptop GPU (8 GB VRAM, Compute 8.9)	All model inference
CPU	Intel Core i9-14900HX (24 cores, 32 threads)	Daemon, frame dispatch, logging
RAM	16 GB DDR5 (15 GB available, 1 GB reserved for iGPU)	System
Camera	Logitech Brio 100	Fixed-focus 1080p USB webcam
Host	Lenovo Legion 7 Gen 9 16IRH9, Linux 6.17, Ubuntu 24.04	Full system
Storage	1 TB NVMe SSD (dual-boot partition; WSL backup instance)	OS, models, logs

## 4.2 Software Stack

Layer	Technology	Version	Role
Vision model	YOLOv8n (Ultralytics)	8.4.63	Fast detection
VLM	Qwen3-VL-4B (INT4, bitsandbytes)	—	Scene understanding
Face detection	MediaPipe BlazeFace	Latest	Identity gate
Daemon	Python + OpenCV	3.11.15 / 4.13.0	Main loop, dispatch
Agent API	OpenAI-compatible HTTP	—	Agent communication
Quantization	bitsandbytes NF4	0.43	Model compression
CUDA runtime	NVIDIA CUDA	12.0	GPU acceleration
Service manager	systemd -user	—	Daemon lifecycle

## 4.3 Daemon Architecture

The vision daemon (`vision-daemon.py`, ~26 KB, 540 lines) runs as a `systemd` user service. The main loop executes at the frame interval (~33 ms for 30 fps capture):

loop:

1. Capture frame from `/dev/video0` (OpenCV)
2. Run YOLOv8n → dorsal results
3. Run MediaPipe → face detection + ID gate
4. Run frame differencing → motion level
5. Classify activity from bbox history
6. If 30s since last Qwen inference:
  - Run Qwen → ventral results
7. Merge dorsal + ventral + motion + identification into observation
8. Write observation to rolling JSON log
9. Advance state machine ticks (cooldowns, timers)

esac

The daemon writes observations to two locations: - `latest.json` — current observation (read by agent integration) - `history.jsonl` — append-only log (for analysis and debugging)

## 4.4 Model Quantization

Qwen3-VL-4B is loaded with 4-bit NormalFloat quantization (bitsandbytes NF4, Dettmers et al., 2023):

```
import torch
from qwen_slow_path import QwenSceneDescriber

describer = QwenSceneDescriber(
    model_id="Qwen/Qwen3-VL-4B-Instruct",
    interval=30.0,
    device="cuda:0",
    load_in_4bit=True,
)
describer.load()
# Result: ~2.8 GB VRAM, 3.3s per inference
```

Without quantization, the same model consumes ~7.6 GB—exceeding the available VRAM alongside YOLO and other processes.

---

## 5. Agent Integration

### 5.1 Pre-LLM-Context Injection

The visual observations are injected into the agent’s context before every conversation turn. The agent platform reads `latest.json` and prepends a formatted sensory preamble to the system prompt:

```
[Senses @ 14:23:04 – You see the following:
  Objects: person, monitor, cup, bookshelf, plant
  Activity: Rob is typing at his desk, facing the screen
  Scene: Rob in a red jacket with glasses and a beard, appears focused and tired
  Motion: Low]
```

This ensures the agent is spatially and semantically aware of the current environment before responding. The injection is transparent—the agent never needs to poll or request visual information; it is always available in context.

### 5.2 Trigger Events (Reflex Arc)

In addition to context injection, the pipeline generates autonomous trigger events when specific visual conditions are met:

Event	Condition	Cooldown	Purpose
Rob returned	Absent $\geq 5$ min, now detected	300 s	Greeting
Rob woke room	First detection, 5-10 AM	3600 s	Morning greeting
Motion detected	Motion level > threshold	30 s	Awareness
Scene changed	Ventral scene diff detected	120 s	Context update

Each event fires through an HTTP API call to the agent platform, which generates an agent response—the agent decides what to say (if anything), maintaining agency over the final output. The API is OpenAI-compatible, allowing the pipeline to interface with any compatible agent runtime.

### 5.3 Privacy Architecture

Principle	Implementation
All local	Zero network egress for raw frames
No recording	Rolling ring buffer, not persistent video
Ephemeral features	Face encodings in memory only
API protection	Bearer token authentication
No cloud dependency	All models run locally

## 6. Evaluation

### 6.1 Performance Metrics

Metric	Dorsal Stream	Ventral Stream
Inference time	16 ms (55 fps)	3,300 ms (0.3 fps)
VRAM usage	~1.0 GB	~2.8 GB
GPU util.	~15-20%	~80-90% (during inference)
Latency to agent	~33 ms (frame interval)	~30 s (sampling interval)

Total VRAM: ~5.6 GB (combined with face detection and embeddings) Free VRAM: ~2.0 GB (on 8 GB RTX 4060)

### 6.2 Continuous Operation

The pipeline has been running continuously for 48+ hours (as of this writing). During this period:

- **2,290,000+ frames** processed (55 fps average)
- **5,760 Qwen inferences** completed (30 s interval)
- **Zero crashes** due to model failure
- **Memory stable** — no VRAM growth detected after initial model loading

### 6.3 Qualitative Observations

The ventral stream produces contextually rich scene descriptions that capture not just objects but *state* and *affect*:

*“Rob in a red jacket, chin resting on hand, wearing glasses and a beard. Expression appears thoughtful but tired — late evening, monitor glow on face. Coffee cup present on the right.”*

*“A dim room with cool blue light from a distant window. No people visible. The space feels quiet and empty.”*

The dorsal stream reliably detects presence/absence transitions and provides continuous spatial awareness. The ventral stream provides the semantic texture required for contextually appropriate agent responses.

### 6.4 Failure Modes

Failure	Cause	Mitigation
False motion detection	Lighting changes (sunlight, shadows)	Adaptive threshold
VLM hallucination	Low-light conditions	Confidence gating
Face ID drift	Extreme angles	Multi-frame consensus
GPU OOM	Concurrent agent inference	VRAM budget monitoring

## 7. Conclusion and Future Work

We have presented a production vision pipeline that implements biomimetic dual-stream temporal asymmetry for a persistent autonomous agent. The 1800:1 ratio between dorsal and ventral stream processing rates demonstrates that extreme temporal asymmetry is not only viable but computationally optimal for agents requiring both continuous spatial awareness and occasional deep understanding.

### 7.1 Implications

This work suggests that the Two-Streams Hypothesis—long a cornerstone of cognitive neuroscience—has direct engineering applications in autonomous systems. The computational efficiency argument (fast and cheap on most frames, slow and expensive on key frames) mirrors the brain’s own resource allocation strategy and may generalize to other sensory modalities.

### 7.2 Future Work

- **SLAM integration:** Adding depth sensing (Kinect v1) for spatial mapping and navigation

- **Hardware platform:** Migrating to a ROS-based bipedal robot (HiWonder Ainex) for embodied vision
  - **Multi-person awareness:** Identity-tiered access control with personality gating per detected person
  - **Top-down attention:** Agent-directed visual attention—choosing what to look at, not just processing everything
  - **Verification loops:** Visual confirmation that agent-initiated actions have succeeded
- 

## References

1. Ungerleider, L.G. & Mishkin, M. (1982). *Two cortical visual systems*. In Analysis of Visual Behavior.
  2. Goodale, M.A. & Milner, A.D. (1992). *Separate visual pathways for perception and action*. Trends in Neurosciences.
  3. Redmon, J. et al. (2016). *You Only Look Once: Unified, Real-Time Object Detection*. CVPR.
  4. Bai, S. et al. (2025). *Qwen3-VL Technical Report*. arXiv:2511.21631.
  5. Dettmers, T. et al. (2023). *QLoRA: Efficient Finetuning of Quantized Language Models*. NeurIPS.
  6. Carion, N. et al. (2020). *End-to-End Object Detection with Transformers*. ECCV.
  7. Liu, H. et al. (2023). *Visual Instruction Tuning*. NeurIPS (Oral).
  8. VanRullen, R. & Thorpe, S.J. (2001). *The time course of visual processing: from early perception to decision-making*. Journal of Cognitive Neuroscience.
  9. Jocher, G. et al. (2020-2023). *Ultralytics YOLO*. GitHub.
  10. Lugaresi, C. et al. (2019). *MediaPipe: A Framework for Building Perception Pipelines*. CVPR Workshop.
- 

Correspondence: [stay.curious@exileresearch.ca](mailto:stay.curious@exileresearch.ca)