‘Meaning Modulates Physics’ Technical Appendix

1. SHR Generation Pipeline: Architecture and Implementation

The SHR data used in this work is generated via a production-ready cloud pipeline. At a high level, the system combines GPU-accelerated feature extraction with a staged reasoning architecture to transform unstructured multimedia inputs into the dual-stream SHR format.

Stage	Haptica Component	Function
Input	Video / Audio	Source multimedia content.
Pre-processing	Multimodal Feature Extraction Layer	Generates a time-coded vector of low-level features (e.g., transcription, acoustic events, visual concepts, structural transitions)
Reasoning	Staged Multi-Pass Architecture	Synthesizes the feature vector into a structured, physically-grounded SHR object
Output	Dual-Stream SHR (a_{phys} + a_{perc})	Final machine-readable format for AI training and human-facing applications

Performance Metrics:

- **Processing Latency:** Present implementation requires several times longer than the input duration to process, with ongoing work to reduce this gap toward real-time performance.
- **Computational Requirements:** Standard datacenter-class GPUs.

---

2. Comparative Study: Physically-Grounded System Output

The following is the direct, uncurated SHR output from our framework for the 8-second AI-generated video. It demonstrates the system's foundational ability to produce a grounded report of a flawed physical reality.

Haptica (Working Blind) — Direct System Output:


```
{
  "beat_type": "The Key in the Cup",
  "events": [
    {
      "event_type": "Transient.Impact.Sharp",
      "timestamp_ms": 3800,
      "duration_ms": 150,
      "physics_descriptor": {
        "intensity": 0.6
      },
      "perceptual_descriptor": {
        "intensity": 0.6,
        "narrative_rationale": "Captures the distinct, high-frequency
clink of the metal key hitting the inside of the ceramic mug..."
      }
    },
    {
      "event_type": "Texture.Impact.Soft",
      "timestamp_ms": 4800,
      "duration_ms": 150,
      "physics_descriptor": {
        "intensity": 0.3
      },
      "perceptual_descriptor": {
        "intensity": 0.3,
        "narrative_rationale": "Represents the soft sound and feeling of
the mug being lifted off the wooden desk."
      }
    }
  ],
  "schema_version": "2.0_unified"
}
```


Analysis: Unlike a baseline LMM, which hallucinates a "magic trick" narrative to resolve the video's physical incoherence, our framework produces a factual, physically-grounded report of the events that occurred. The system's perception is anchored in the verifiable physics of the audio cue, demonstrating true grounded comprehension.

---

3. Mathematical Formalization

The SHR framework is built on a dual-stream principle designed to separate objective physics from subjective context.

- Physics Stream (Context-Invariant): Provides a set of objective physical descriptors intended to approximate an event's ground-truth dynamics.
  $a_{\text{phys}} = g_\theta(x)$
  (where x is the set of input features and g is a model with parameters θ)
- Perceptual Stream (Context-Modulated): Provides a narrative-aware interpretation of the physics.
  $a_{\text{perc}} = m_\varphi(a_{\text{phys}}, c)$
  (where c is the context vector)
- Modulation Mechanism (FiLM form): A preferred embodiment for the modulation function m_φ is Feature-wise Linear Modulation, where the context vector c learns to scale (γ) and shift (β) the physics stream feature-wise.
  $a_{\text{perc}} = \gamma(c) \odot a_{\text{phys}} + \beta(c)$
  (where $\odot$ is element-wise multiplication)

---

4. Foundational Intellectual Property

- **Granted:** U.S. Patent 10,429,931. “Systems for multimedia tactile augmentation” (Issued Oct. 1, 2019).
- **Pending:** “Context-Aware Physical Event Annotation System for Training Physically-Intelligent AI.”
- **Note:** These patents cover methods for synchronized haptic/media delivery and for generating physically-grounded, context-aware data. Any derivative implementation of this architecture may overlap with these claims.

---

5. Project Status and Contact

The pipeline and SHR data format described here were architected and built by the author. The pipeline has been validated at pilot scale and successfully used to generate the data presented in this paper. Future work will extend validation across larger datasets.

For technical inquiries or collaboration, please contact the author at:
alan [at] haptica [dot] io