FRENETIC

MEASURING WHAT SCREENS DO TO US

An open standard for measuring the sensory load of audiovisual content, and a personalised way to bring it within each viewer's tolerance.

From signal to comfort

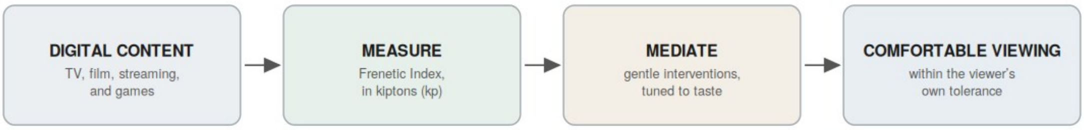

The signal is measured, then brought within the viewer's own tolerance.

White Paper

Version 1.0, 2026

Srinivas Kasturi

Contents

Executive summary

Audiovisual content has become cognitively expensive. Modern television, news, streaming and games are cut faster, layered more densely, and lit more aggressively than the material of a generation ago. For a large group of viewers — autistic people, those with attention or sensory-processing differences, migraine-prone and post-traumatic viewers, and many older audiences — this can make ordinary content tiring, distressing, or impossible to watch. Subtitles do not fix it. Audio description does not fix it. Turning the volume down barely touches it.

Frenetic introduces a way to measure the problem and a way to relieve it. It defines the Frenetic Index, a measurement of the cognitive-perceptual sensory load that audiovisual content places on a viewer, expressed in a unit called the kipton (symbol kp) on a bounded scale from 0 to 100. The index is built from four orthogonal components of the signal and weighted to each individual through a short calibration. An interoperable metadata format lets content carry its own load data, the way it already carries subtitles or high-dynamic-range information; where no such data exists, the load is measured on the fly.

A viewer calibrates once to find their personal comfortable threshold. When content rises above that threshold, an intervention engine applies the lightest available change — softening a hard cut, calming a shaky shot, easing a flash, lifting a voice clear of layered sound — just enough to bring the work within tolerance without sanitising it. The same standard works in an inline adaptor between a player and a screen, inside a television, inside a streaming service, or as software on a general-purpose device. The methods described here are the subject of United Kingdom patent application GB2611710.1, filed by the inventor, Srinivas Kasturi. The aim is a platform and a standard, not a single product.

1. The problem

Content that cannot be watched. A significant share of the population is more sensitive to sensory input than the median viewer that mainstream content is made for. Fast editing, restless camerawork, dense sound design and bright flashing can push these viewers past the point of comfort and into distress, fatigue, or withdrawal. The result is exclusion from culture and from shared information, including news.

Loudness is not load. Broadcast loudness standards measure how loud audio is and even out the peaks, which protects hearing and keeps programmes consistent. They say nothing about cognitive load. A scene can sit within every loudness limit and still be acoustically chaotic — many layered sources, dense transients, competing speech — in a way that overwhelms a sensitive listener. Load is not the same quantity as loudness, and no existing standard measures it.

Accessibility that skips the senses. Captions substitute text for audio; audio description narrates the picture. Both are valuable, and neither addresses the sensory intensity of the content itself. There is no accessibility track that describes, or moderates, how much the work demands of the viewer's nervous system.

One setting for everyone. Where mitigations exist at all, they tend to be a single switch — a blanket "reduce flashing" toggle that is either on or off, identical for every person. Sensitivity is not uniform: it differs between people, between conditions, and within one person from day to day. A binary, one-size setting cannot meet it.

Closed and per-platform. What little exists is locked to a particular operating system or a particular app, applied only to that platform's own rendered video. There is no shared, interoperable way for a measurement to travel with content across devices and services, and no way for a viewer's personal needs to follow them from one screen to the next.

The pattern is consistent: sensory load is real, individual, and consequential, yet it is unmeasured, unlabelled, addressed only by crude all-or-nothing switches, and never in a form that crosses platforms.

2. The solution

Frenetic answers each of these problems with one coherent design: a measurement, a way to carry it, a way to personalise it, and a way to act on it gently.

2.1 A unit for sensory load

The Frenetic Index expresses the sensory load of a moment of content as a number of kiptons, from 0 to 100, divided into five readable bands: Calm, Steady, Active, Intense, and Frenetic. The scale is bounded, so the most demanding content approaches but never exceeds 100, and every viewer's comfortable limit can be placed at a single point along it.

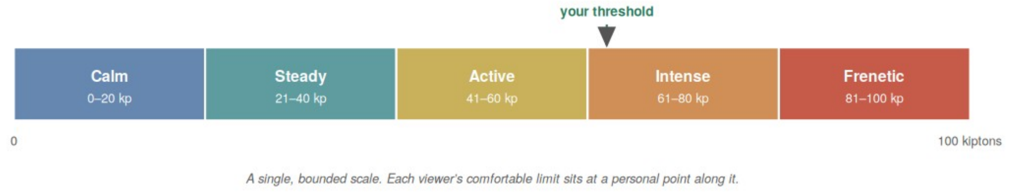

Figure 1. The kipton scale and its five bands. A viewer's comfortable threshold sits at a personal point along it.

2.2 Four components

The index is composed from four orthogonal measurements of the signal: temporal velocity, how fast the content cuts and moves in time; spatial chaos, how busy and disordered each frame is; acoustic density, how layered and transient-dense the sound is, independent of loudness; and luminance volatility, how sharply and often brightness changes. Keeping the components separate means the cause of a high reading is always legible, and that interventions can target the offending dimension precisely.

2.3 A composition that compounds

The four components are combined into a single value by a composition function under each viewer's personal weights. The preferred form is multiplicative and bounded, which captures an important truth about perception: load compounds across dimensions. A busy picture and a dense soundtrack together are harder to bear than either alone, and the composite reflects that, while remaining within the 0-to-100 scale.

2.4 Personal calibration

Each viewer calibrates once, through a short setup that presents controlled examples and records where comfort gives way to strain. From this the system derives a personal threshold and the personal weights that say which dimensions matter most for that individual. A person highly sensitive to motion but tolerant of sound will be calibrated differently from someone for whom the reverse is true. The result is a profile that belongs to the viewer.

2.5 An interoperable metadata format

Frenetic defines a metadata format so that content can carry its own load data alongside the picture and sound, announced and transported by the same mechanisms that already carry subtitles and high-dynamic-range information. Because it rides on accepted infrastructure, it coexists with captions, audio description, signed-language tracks and colour metadata rather than competing with them, and devices that do not understand it simply pass it by.

2.6 Inference where there is no metadata

Most content today carries no load data. For this case the playback device measures the four components itself, from the signal, using a short rolling buffer that lets it see slightly ahead and prepare a smooth response. A viewer therefore benefits immediately, on existing content, while authored metadata improves accuracy over time as it becomes available.

2.7 The lightest intervention

When content rises above the viewer's threshold, an intervention engine applies the gentlest change that brings it back within tolerance. Hard cuts are softened into brief cross-dissolves; restless handheld motion is digitally steadied; the periphery of a busy frame is calmed while the centre stays sharp; sudden flashes are eased; layered audio is reduced so dialogue stays clear. The depth of each intervention scales with how far the content has exceeded the threshold, and authorial markers let creators protect moments that are meant to be intense. The goal is always to preserve the work, not to flatten it.

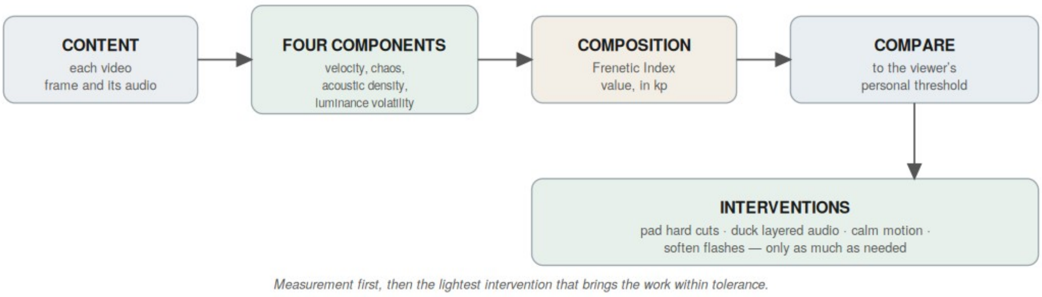

Figure 2. Measurement first, then the lightest intervention that brings the work within tolerance.

2.8 Learning from the viewer

A viewer can always signal that the system is doing too much or too little for the content in front of them. Frenetic learns from these signals, adjusting a contextual offset to the threshold for similar material, so that the experience converges on the individual's real preferences rather than a fixed assumption.

3. Under the hood

3.1 The measurement pipeline

Each of the four components is computed from the decoded signal. Temporal velocity draws on shot-change detection and motion estimation; spatial chaos on spatial-frequency and edge-density analysis; acoustic density on source separation and transient and spectral measures of the audio; luminance volatility on per-region brightness change over time. The components are normalised to a common range before composition.

3.2 The composition function

Under the preferred multiplicative form, each normalised component reduces the remaining "comfort space" in proportion to its personal weight, and the composite is the total reduction expressed on the 0-to-100 scale. This guarantees boundedness and reproduces the compounding of load across dimensions. A simpler weighted-sum form is also defined for constrained devices.

3.3 Metadata and carriage

Load data is carried per frame within the content's existing user-data channels, with its presence announced in the programme's table of contents, mirroring the way other accessibility and colour metadata are signalled. Authorial markers and a small set of per-moment flags travel in the same payload. The format is designed to be checked by conformance tools built on existing stream analysers.

3.4 The intervention engine

The engine holds a menu of interventions, each addressing one or more components, each with an onset and a depth set by the viewer's preferences and by the current overshoot. Interventions ramp in and out smoothly to avoid drawing attention to themselves, and defer to authorial markers that protect intended moments.

3.5 Where it runs

The same measurement, comparison and intervention can be performed in an inline adaptor placed between a source and a display, inside a television or monitor, inside a set-top box, streaming stick or console, on the servers of a streaming service before delivery, or as software on a general-purpose device. The standard is one thing; the places it can run are many.

4. The ecosystem

Frenetic is most valuable as a shared standard that many parties can adopt independently. A single device helps one viewer; an interoperable standard helps the whole audience.

Figure 3. Viewers, content producers, device makers and services around one open measurement standard.

4.1 Roles

- **Viewers** calibrate once and carry a personal profile that follows them across screens.

- **Content producers** author load data into their works during post-production, live broadcast, or game rendering, and protect intended moments with authorial markers.
- **Device makers** implement measurement and mediation in displays, source devices and adaptors.
- **Streaming services** carry the metadata and can offer Frenetic as an accessibility feature alongside captions and audio description.

4.2 The open standard and metadata

The measurement, the unit, the composition function and the metadata format are specified openly so that any party can implement them and conformance can be tested. This is what lets load data travel with content from producer to viewer regardless of which devices and services lie between.

4.3 Calibrate once, carry everywhere

Because the profile belongs to the viewer rather than to a device, the same calibration applies on a living-room television, a tablet on a train, or a friend's screen, given a Frenetic-aware device or app. Personalisation follows the person, not the hardware.

4.4 How value is created

Value here is access, not a toll. Viewers gain content they could not previously watch; producers and services reach an audience they were previously losing and meet accessibility expectations; device makers gain a differentiating feature. The standard is silent on commercial terms, which are left to those who implement it.

> *A note on scope. This paper describes intended design and direction. It is not a commitment to deliver any specific feature, nor an offer of any product, security or token, nor legal or financial advice. Details may change as the standard develops.*

5. In practice

The viewer. Maya, an autistic adult who had stopped watching the evening news because the pace left her exhausted, calibrates once on her television. The bulletin now plays with its hardest cuts softened and its layered studio sound eased, comfortably within her tolerance, and she follows the day's events again.

The parent. A father calibrates on behalf of his young son, who cannot reliably report his own limits, and sets a gentle profile for the cartoons the boy loves but that used to end in meltdowns.

The broadcaster. A public broadcaster adds a Frenetic authoring stage to its live news workflow, so that load data is computed and embedded in real time and protected moments are flagged by the producer, reaching every Frenetic-aware screen downstream.

The device maker. A television manufacturer implements the standard in its accessibility settings, so that any source reaching the panel — broadcast, streaming app, or games console — can be brought within a viewer's tolerance without extra hardware.

6. Design principles

- **Measure before you mediate:** a transparent, bounded measurement comes first, so that any change is principled and legible.

- **Personal, not uniform:** a calibrated threshold and weights per viewer, not a single switch for everyone.
- **Preserve the work:** the lightest intervention that suffices, with authorial markers protecting intended intensity.
- **Interoperable by design:** an open unit and format that travel with content across devices and services.
- **Honest about the landscape:** flashing-reduction toggles and loudness standards already exist; the contribution here is a personalised, cross-modal, interoperable measurement of load and a graduated way to act on it.

7. How this differs

The table below sets common approaches today against the design described here. It aims to be fair rather than complete, and the features of other products vary.

Dimension	Common approaches today	Frenetic
What is measured	Loudness, or nothing	Cognitive-perceptual load, in kiptons
Modalities	Single dimension (e.g. flashing)	Four components across picture and sound
Personalisation	One setting for everyone	Calibrated threshold and weights per viewer
Action	On/off, or none	Graduated, work-preserving interventions
Portability	Locked to one platform	Open metadata that travels with content
Authoring	Not authored	Producers embed load data and protect moments

8. Intellectual property and open source

The methods described in this paper — the measurement of audiovisual sensory load on a bounded personal scale, the interoperable metadata format that carries it, the personal calibration that sets each viewer's threshold and weights, and the graduated intervention engine — are the subject of United Kingdom patent application GB2611710.1, filed by Srinivas Kasturi. In keeping with the project's intent that Frenetic become a shared standard, reference tooling and conformance materials are intended to be released openly, so that the community can implement, audit and extend the platform while the core inventions stay protected.

9. Roadmap

Phase 1, Foundations. The measurement, the kipton scale, personal calibration, the inference-mode pipeline, and a first single-device experience.

Phase 2, Carriage and authoring. The metadata format, authoring tools for post-production, live broadcast and game engines, and conformance materials.

Phase 3, Adoption. Implementations across displays, source devices and streaming services, and the portable profile that follows the viewer.

Dates and scope are indicative and will be confirmed as the project develops.

10. Conclusion

Sensory load is real, individual, and until now unmeasured. By naming it, putting it on a bounded personal scale in kiptons, letting content carry it, and acting on it with the lightest touch that brings a work within a viewer's tolerance — all in an open, interoperable form — Frenetic sets out to bring back into reach the screens that had become too much to watch.

Appendix A: Glossary

Frenetic Index. A measurement of the cognitive-perceptual sensory load of audiovisual content, on a bounded scale from 0 to 100.

Kipton (kp). The unit of the Frenetic Index.

Components. The four orthogonal measurements — temporal velocity, spatial chaos, acoustic density and luminance volatility — from which the index is composed.

Personal threshold. The kipton value above which a particular viewer finds content uncomfortable, found by calibration.

Metadata-aware and inference modes. Reading load data carried with the content, or measuring it from the signal where none is carried.

Intervention. A real-time, graduated change to the content that reduces load while preserving the work.

Appendix B: Important notices

This document is for information only and describes intended design and direction. It is not an offer or solicitation to buy or sell any security, token or other instrument, nor a promise of future features, availability or returns, nor legal, financial or investment advice. Statements about future plans are forward-looking and may change. Trademarks and product names of third parties belong to their respective owners and appear here only for comparison. The methods described are the subject of a pending patent application, and nothing here grants any licence.