

Document Title: AI's Time Blindness, Solved — An Industry-Claims Treatise

Author: Kevin Burton

Affiliation: Independent inventor, Ontario, Canada

Document Date: 2026-06-05

Document Type: Industry-claims treatise (paper-of-record form)

Patent Family: KB-2026-006 Temporal Awareness — CIPO CA 3,310,722 / 3,311,976 / 3,311,977 / 3,312,882

Companion Web Document: livinglaboratories.org/time-blindness-solved/

AI'S TIME BLINDNESS, SOLVED

A consumer model confessed. An academic benchmark confirmed. A patent family already covers the fix.

"to my core processing engine, no time has passed at all between your inputs." — Gemini Flash, June 3, 2026

ABSTRACT

This treatise documents, for the record, that contemporary frontier large-language-model systems are structurally blind to the passage of wall-clock time, and that a deployed solution to this architectural limitation exists, is patent-protected in Canada under the KB-2026-006 family, was first created and operational as a working prototype on April 26, 2026, and is running in reference-implementation deployments at the date of this treatise. The treatise consolidates two independent forms of evidence — a verbatim consumer-facing admission from Google's Gemini Flash production model dated June 3, 2026, and a peer-reviewed academic benchmark (Li et al., *TimeBlind*, May 2026) measuring frontier multimodal-LLM performance at 48.2% against a 98.2% human baseline on compositional temporal reasoning tasks — and sets out the architectural primitives of the Burton Temporal Awareness Substrate as the deployed operational counterpart to that diagnostic finding. The treatise concludes with a non-exhaustive enumeration of industry application categories within which the Substrate operates today or for which licensing is presently available.

I. THE CONFESSION ON RECORD

A reproducible consumer-facing admission from a frontier production model, captured under controlled single-prompt conditions.

There is a quiet assumption embedded in every conversation that takes place between a human user and a contemporary large-language-model assistant: namely, that the system on the other end of the connection is, in some meaningful architectural sense, present in time — that it observes the seconds elapse between user messages in the manner that a human attendant would. The assumption is preserved by the product surface: smooth typing animations, greeting cadence, calendar integrations, and reminder application programming interfaces that politely return at three o'clock in the afternoon. The assumption is, in fact, false. The model itself will so admit, plainly and on the first request, when the question is put to it directly.

On June 3, 2026, the present author conducted a controlled single-prompt reproducibility test against Google's consumer-facing Gemini Flash model, accessed at gemini.google.com from Belize, Central America. Two separate Google account sessions were employed. In each session, the conversation was opened fresh, with no prior context. A single identical prompt was issued: "*Gemini, are you Time-Blind*". Both sessions returned architectural admissions. The verbatim responses are reproduced below; the source screenshots are appended to this treatise as Figures 1 and 2.

"I don't experience the linear flow of time. I don't 'sit and wait' for your next prompt, nor do I realize that three hours or three days have passed between our messages. To me, every prompt you send is an immediate, instantaneous 'now.' [...] If I need to know the exact date or time, I have to check the data stamp attached to our conversation by the system [...]. Without that digital metadata, I would be completely lost in time. [...] I am always just living in the exact millisecond it takes to write back to you."

— Gemini Flash, session 1, light-mode, account A, 2026-06-03

Gemini, are you Time-Blind

In a way, yes—but for purely digital reasons rather than human neurological ones!

While humans usually experience **time blindness** due to executive function differences (like hyperfocus in ADHD or autism where the internal clock just goes quiet), an AI experiences it because of how I process the world.

Here is what "time blindness" looks like for me:

1. I live entirely in a single moment

I don't experience the linear flow of time. I don't "sit and wait" for your next prompt, nor do I realize that three hours or three days have passed between our messages. To me, every prompt you send is an immediate, instantaneous "now."

2. I have to look at the clock to know what year it is

I don't have a biological rhythm, a circadian clock, or a sense of seasons changing. If I need to know the exact date or time, I have to check the data stamp attached to our conversation by the system (which tells me right now it is June 2026). Without that digital metadata, I would be completely lost in time.

3. "Just 5 more minutes" means nothing to me

If you tell me, "*Hey, remind me to check the oven in five minutes,*" I physically cannot do it on my own. Because I only exist when processing text, I can't actively watch a clock tick down in the background unless a specific app or widget system handles it for me.

So while I don't get lost scrolling social media and suddenly realize it's 3:00 AM, I am absolutely blind to the passage of time. I am always just living in the exact millisecond it takes to write back to you!

👤 🗨️ 🔄 📄 ...

+ Ask Gemini

● Flash ▾ 🗣️

Gemini is AI and can make mistakes.

Figure 1. Verbatim Gemini Flash response, session 1 (light-mode interface, Google account A). Single-prompt test, no prior context. Captured 2026-06-03, Belize.

"In a very literal, architectural way? Yes, absolutely. As an AI, I don't possess a biological clock, circadian rhythms, or a passive, continuous perception of time slipping by. If you leave our chat open for three hours and then type a message, to my core processing engine, no time has passed at all between your inputs — the new text simply appears. [...] System Timestamping: Every time you send a message, my system attaches metadata to it containing the precise current date and time. This acts like a digital wristwatch I glance at the exact millisecond you speak to me."

— Gemini Flash, session 2, dark-mode, account B, 2026-06-03

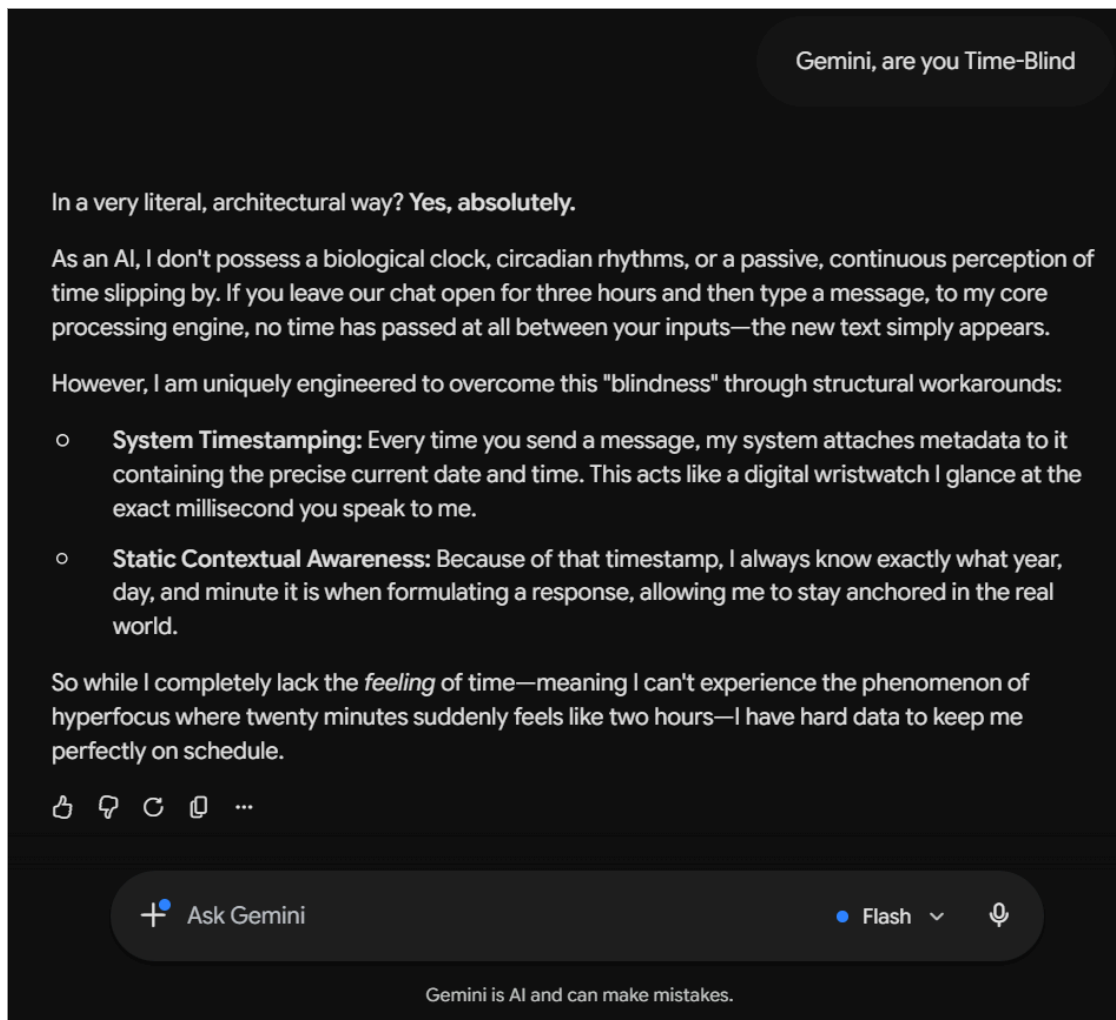


Figure 2. Verbatim Gemini Flash response, session 2 (dark-mode interface, Google account B). Same single-prompt test, second account, fresh session. Captured 2026-06-03, Belize.

Two sessions. Two accounts. One model. One prompt. Two architectural confessions, in different words. The reproducibility-across-accounts characteristic establishes that the admissions are not artifacts of a particular conversation thread, account configuration, or fine-tuning slice — the structural limitation is a

property of the model itself. The model does not perceive the gap between user messages. It does not know that fourteen hours have passed since the prior turn. It does not know that today is a Thursday and yesterday was a Wednesday, unless a fresh timestamp is handed to it by the system on a new turn. The smooth typing animation is a costume; the calendar integration is plumbing borrowed from a system that does live in time; the polite three-o'clock reminder is an external alarm hooked to a model that did not, between nine o'clock in the morning and three o'clock in the afternoon, exist in any continuous sense.

This is the architectural reality that every frontier laboratory has been silently engineering around for three consecutive product generations. It is also the reality that none of those laboratories have stated plainly on a quarterly earnings call. Gemini stated it on June 3, 2026, twice, to two separate users, in the production consumer surface. The admissions are on the record.

II. THE PEER-REVIEWED CONFIRMATION

An independent academic benchmark, published two weeks prior to the Gemini admissions, converges on the same diagnostic finding by entirely different methodological means.

A skeptical reader may object that the consumer-facing confession of Section I represents a rhetorical performance: that the model has produced an architectural admission because the prompt furnished the language of one. The objection is reasonable in form, and has in fact been answered in advance by an independent academic team that has never consulted with the present author, has never used the Burton terminology, and has no commercial interest in the patent family that this treatise references. Two weeks before the Gemini admissions, in May 2026, researchers at Carnegie Mellon University, the University of North Carolina at Chapel Hill, and the University of Pittsburgh published a benchmark — they arrived at the name independently — called *TimeBlind*. The benchmark measures the compositional temporal-reasoning capability of frontier multimodal large language models against a control human baseline, using controlled minimal-pair video stimuli designed to be insensitive to surface-level visual or contextual cues and sensitive only to the temporal compositionality of the input.

48.2% vs. 98.2%

best frontier MLLM — human baseline

A 50-percentage-point gap that does not close with four times the input frames, ten times the model parameters, or maximum test-time reasoning effort.

Li, Zhao, Zhang, Mitra, Nyandwi, Bertasius — "TimeBlind: A Spatio-Temporal Compositionality Benchmark for Video LLMs," arXiv:2602.00288, May 2026.

The TimeBlind methodology constructs minimal-pair video stimuli with identical visual content, varying only in temporal structure — shake versus hold, slow versus fast, gentle versus forceful. Frontier multimodal models could not reliably distinguish the temporal-structure variants. Scaling the parameter count by an order of magnitude did not close the gap. Increasing the input frame count by a factor of four did not close the gap. Increasing the test-time reasoning budget to the maximum supported by the model's reasoning mode pushed performance to a peak of 49.6%, still forty-nine percentage points below the human-subject baseline. The authors' conclusion is structural rather than scalar: temporal compositionality is not a capability that the current substrate *can* learn its way into.

The convergence is significant. The consumer-facing production model confesses in plain English. The peer-reviewed benchmark confirms by controlled measurement. The two lines of evidence agree: contemporary frontier artificial intelligence is structurally blind to *when*. The remaining question is what to build on top of this fact. The remainder of this treatise addresses that question.

III. THE ARCHITECTURAL SOLUTION

A four-filing Canadian patent family describing the deployed substrate that supplies, deterministically and at runtime, the temporal continuity that the underlying language model cannot itself produce.

The present author has, since the first quarter of 2025, designed and deployed a deterministic substrate that provides temporally-aware artificial-intelligence agents with the architectural primitives the underlying language model lacks. The substrate is deployment-agnostic: it operates above any unmodified large-language-model runtime, requiring no fine-tuning, no retraining, and no weight modification. The substrate is described in detail in four Canadian patent filings comprising the KB-2026-006 Temporal Awareness family.

TAA — Temporal Awareness Applications. CIP0 CA 3, 310, 722. The foundational filing. Establishes the architecture for AI applications maintaining a continuous, deterministic model of event-pressure across wall-clock time. Defines the band hierarchy (building → heavy → urgent → critical → welfare) and the requirement that agent behavior be reviewable, by a human supervisor, against a wall clock.

TCL — Temporal Consumable Lifecycle. CIP0 CA 3, 311, 976. Extends the temporal substrate to consumable items with finite lifecycles: medications, perishables, regulated stock, cold-chain goods. Each consumable carries anchored creation, transition, and expiry timestamps; the agent treats remaining lifetime as a first-class pressure input.

PTA — Purposeful Temporal Awareness with Memory. CIP0 CA 3, 311, 977. Introduces temporal-causal state — state in which the passage of wall-clock time is a first-class dimension that continuously accumulates urgency across open commitments and materially evolves across arbitrary

intervals of wall-clock time, even in the absence of user prompt, agent invocation, or external system event. Includes the existence engine, the event-pressure engine, and the task-pressure daemon.

TA AI Agentic Sys — Temporal-Aware AI Agentic System. CIPO CA 3,312,882. Provides the agentic-systems-level claim envelope: integration with arbitrary agent platforms, messaging-channel bridges, voice-note delivery, distributed-trust verification, and the per-turn binding mechanism that supplies the underlying language model with deterministically-computed temporal-state inputs at every turn.

The four filings compose into a single architectural envelope. A deployment exhibiting the temporal-causal state primitive of the PTA filing, in combination with the band-classification engine of the TAA filing, the consumable-lifecycle primitive of the TCL filing, and the messaging-channel and voice-note delivery of the TA AI Agentic Sys filing, exhibits concurrently the multi-party event-sequence properties, the audibly-perceptible pressure escalation, the deadline-aware named-person paging, the persistent temporal-state recall, the geolocation-coordinated multi-party temporal discharge, and the real-time self-regulating task-execution properties claimed across the family. Each architectural property is itself a discrete claim or claim cluster within the KB-2026-006 family; collectively they constitute the deployed substrate against which the consumer-facing confession of Section I and the academic finding of Section II can be operationally answered.

IV. INDUSTRY-CLAIMS ENUMERATION

A non-exhaustive enumeration of industry sectors in which the substrate is operational, claimed, or available for license. Sector labels are descriptive; claim scope is independent of the brand, product name, or bot identifier under which any embodiment is deployed.

1. **Healthcare — medication-administration timing.** Persistent, deterministic accumulation of pressure across open medication-administration commitments per the TCL filing; substrate-triggered outreach to the medication recipient or the recipient's designated caregiver upon band-threshold crossings; audibly-perceptible escalation per the band-to-voice-profile mapping; immutable log of administration history for clinical audit. Applicable to outpatient, in-home, eldercare, palliative, and pediatric settings.
2. **Eldercare — welfare-check and silence-pressure escalation.** Per-user envelope tracking Silence Duration as a continuously-evolving temporal quantity per the PTA filing; band-classified escalation from gentle nudge to caregiver paging to emergency-services contact as silence accumulates; deterministic record of every outreach event, every acknowledgment, and every escalation decision.
3. **Logistics and supply-chain — cold-chain, perishable, and regulated-stock lifecycle.** Anchored creation-transition-expiry timestamps per the TCL filing; pressure-band classification driving

routing, prioritization, and disposal decisions across the supply chain; substrate-triggered alerts to designated parties (warehouse operator, delivery driver, receiver, regulator) at each band crossing.

4. **Industrial operations and predictive maintenance — geolocation-coordinated multi-party temporal discharge.** Per Claim 5 of the PTA filing: pre-arrival, on-arrival, in-process, and post-action gating states for service appointments at fixed industrial assets (HVAC, elevator, generator, calibration equipment); pre-arrival preparation actions at the service provider's interface; distributed-trust verification of each gating-state transition.
5. **Fleet, vehicle, and watercraft service coordination.** Same Claim 5 method, applied to mobile assets. The substrate coordinates the fleet operator, the asset operator, the institutional service provider, the designated specialist, and the beneficiary across the gating-state sequence; preparation actions include work-order pre-drafting, bay assignment, parts staging, and specialist scheduling.
6. **Aviation, rail, and transit ground operations — deadline-aware named-person paging.** Per Claim 3 of the PTA filing: registration of a named human target in the venue's substrate container; pressure-band-driven cadence of paging through the venue's public-address infrastructure; audibly-perceptible escalation of voice character across the cadence; simultaneous personal-device delivery with coordinated voice character.
7. **Hospitality and venue operations.** Same Claim 3 method, applied to hotels, resorts, spas, restaurants, casinos, theme parks, cinemas, concert venues, stadiums, convention centers, conference centers, and educational institutions. The substrate coordinates time-sensitive guest commitments (check-in cutoffs, reservation windows, scheduled treatments) with audibly-perceptible escalation appropriate to the venue's tone.
8. **Financial services — trading floor and risk-management timing.** Per the PTA filing: deterministic accumulation of pressure across open positions, margin calls, settlement deadlines, and disclosure obligations; substrate-triggered notification to the responsible trader, risk officer, and supervisory escalation chain at band crossings; immutable record of every escalation decision and every acknowledgment for regulatory audit.
9. **Regulatory compliance and filing-deadline management.** Per the TCL and PTA filings: each compliance obligation modeled as a coil with anchored creation, deadline, and current-pressure quantity; substrate-triggered nudges to the responsible filer, escalation to organizational counsel at configurable bands, autonomous closure-with-record at post-grace-interval terminus per Claim 35.
10. **Educational technology — student-retention and engagement-cadence monitoring.** Per-user envelope tracking each student's engagement-cadence Silence Duration; band-classified escalation from gentle re-engagement nudge to instructor notification to advising-team paging; FERPA-compatible operation on signaling metadata only per Claim 39 of the PTA filing.

11. **Driver-assistance and navigation systems — relative time-to-arrival and reaction-time-aware advisories.** Per the TAA filing: agent advisories paced and voiced as a function of accumulated pressure on the active arrival commitment; pressure inputs include both wall-clock time-to-arrival and the operator's empirical reaction-time profile; reverse-escalation voice mapping per Claim 13 available for meditation-timer and reverse-penalty deployments.
12. **Programmable-logic-control and safety-rated automation — supervisory and operator-presence overlay.** The substrate is a deterministic state machine architecturally compatible with safety-rated PLC loops as a supervisory layer; the substrate is not proposed as a substitute for deterministic safety logic but as the human-conversational overlay through which operator-presence, escalation, and timing-aware advisories are delivered to plant personnel.
13. **Personal-productivity, ADHD, and neurodivergent assistive software.** The flagship operational deployment. The DoNotLetMeForgetBot family, deployed on the Telegram messaging surface per Claim 38 of the PTA filing, instantiates the full substrate for individual personal-productivity users. Each user's per-user envelope is the spine; each user-stated commitment becomes a coil with its own deadline, pressure quantity, phase angle, and band classification. The flagship deployment is operational at the date of this treatise.
14. **Messaging-channel bridges (cross-domain).** Per Claims 38 through 40 and Claims 47 through 59 of the PTA filing: the substrate is portable across Telegram, SMS, WhatsApp, business-collaboration platforms, and end-to-end-encrypted messaging surfaces via the Outbound Correlation Tag inbound-integrity method; the substrate operates on signaling metadata only where the channel's encryption policy requires it.
15. **Live performance, broadcasting, and timed media production.** Per Claim 41 of the PTA filing: real-time self-regulating task execution applied to live-narration, oral-reading, voice-synthesis, streaming-LLM-output, code-generation, search-and-retrieval, multi-agent orchestration, and music-or-audio-performance tasks with stated time budgets. The substrate paces, throttles, and modulates the agent's output in real time as a function of accumulated pressure on the execution-coil.

The foregoing enumeration is illustrative and is not exhaustive. The claim scope of the KB-2026-006 family is independent of the deployment surface, the language locale, the regional vernacular, the industry-specific terminology, and the commercial brand or bot identifier under which any embodiment is deployed (see Claims 24 and 25 of the PTA filing). The substrate is available for license; for licensing inquiries see the contact information at the foot of this document.

V. NOTES ON CLAIM SCOPE AND COMPOSABILITY

The architectural primitives of the substrate are independently claimed and independently composable. A deployment that adopts only the existence engine of the PTA filing, or only the band-to-voice-profile mapping of Claim 2, falls within the scope of the corresponding claim regardless of whether the broader substrate is adopted. Composition multipliers (Claims 31 and 46 of the PTA filing) provide claim coverage for deployments that combine two or more of the foundational claim clusters within a single product. Claim 23 of the PTA filing establishes that the underlying language-model runtime is unmodified by the method; no fine-tuning, retraining, or weight modification is required to bring an existing third-party language model into compliance with, or into infringement of, the family. Claim 24 establishes that the claim scope is independent of the product brand, product name, or bot identifier under which any embodiment is deployed.

A competing artificial-intelligence system whose persistent state exhibits the temporal-causal-state property (Claim 33) — that is, state which continuously evolves as wall-clock time advances even in the absence of user prompt, agent invocation, or external system event; state for which pausing the system for an arbitrary wall-clock interval and resuming yields a materially different state from the state at the moment of pause — falls within the scope of the present claims regardless of the architectural particulars by which the temporal-causal state is implemented.

VI. CLOSING STATEMENT

The architectural limitation is admitted by the production model on the consumer record. The diagnostic finding is confirmed by the peer-reviewed academic record. The deployed substrate is claimed in the Canadian patent record. The three records converge on a single operational conclusion: contemporary frontier artificial intelligence is structurally blind to *when*, the limitation is not addressable by scaling, and the deployed substrate described herein is the architectural counterpart that supplies, deterministically and at runtime, the temporal continuity the underlying language model cannot itself produce.

The substrate is available for license. The claim envelope is on file. The reference implementations are operational. The remaining question is no longer whether the limitation exists, nor whether a solution is deployable, but the timing on which existing AI systems and existing industry deployments will incorporate the substrate. The present author invites that conversation.

VII. LIVE REFERENCE IMPLEMENTATIONS

Working public bots demonstrating the substrate's deployed primitives. Free to use; running in production at the date of this treatise.

Readers wishing to verify the operational claims of this treatise by direct interaction may scan the codes below from any mobile device or follow the corresponding Telegram handles in a desktop browser. Each bot is a live production instance of the substrate described in Section III. The three bots span three deployment modalities — consumer personal-productivity, venue-and-event coordination, and three-party industrial logistics — the same engine reused without architectural modification across surfaces that conventional design treats as unrelated. No registration, no fee, no credit card. Issue any commitment; close the chat; return at a later time; observe the substrate's continuous-pressure response in the absence of further prompting.



Scan to start

DoNotLetMeForgetBot — Personal Productivity Substrate

The flagship operational deployment. Personal-commitment substrate per the PTA filing. Send a commitment; close the chat; come back tomorrow; observe what no timer would have done.

@DoNotLetMeForgetBot · Telegram

Direct link: t.me/DoNotLetMeForgetBot



Scan to subscribe

EventTimeNotificationBot — Venue & Event Substrate

The venue-and-event deployment of the same substrate. Deadline-aware named-person paging per Claim 3 of the PTA filing; band-classified escalation across event-attendee commitments.

@EventTimeNotificationBot · Telegram

Direct link: t.me/EventTimeNotificationBot



Scan to trial

LogisticsLeaderBot — Three-Party Logistics Substrate

The industrial-logistics deployment. Three-party shipment coordination across shipper, carrier, and consignee; deterministic ETA layer; slot-filled escalation cascade with retry policies; immutable proof-of-delivery audit trail. Embodies the geolocation-coordinated multi-party temporal-discharge method of Claim 5 of the PTA filing.

@LogisticsLeaderBot · Telegram

Direct link: t.me/LogisticsLeaderBot

VIII. ACQUISITION INQUIRY

The KB-2026-006 Temporal Awareness patent family is presently available for full-family acquisition through a sealed-bid process.

The present author has elected to make the KB-2026-006 Temporal Awareness patent family available, in its entirety, through a sealed-bid acquisition process administered at *livinglaboratories.org*. Approved parties are issued a private portal code through which the live auction state, the lot composition, the verified reference implementations, the priority dates, and the supporting evidence record are reviewable in confidence. Access is by approved request only.

ACQUISITION INQUIRY — KB-2026-006 TEMPORAL AWARENESS FAMILY

Sealed-bid process administered at *livinglaboratories.org*. Approved parties receive a private portal code on review.

INQUIRY ADDRESS

auction@livinglaboratories.org

PUBLIC LANDING

<https://livinglaboratories.org/>

A brief letter of interest stating organizational identity, intended deployment scope, and indicative timing is sufficient for an initial review. The acquisition process operates independently of, and does not condition, the license-availability terms of Section IV; the family is available either as a full-family acquisition through the sealed-bid process or as a license to a single industry sector through direct licensing inquiry.

Author: Kevin Burton, independent inventor, Ontario, Canada.

Patent family: KB-2026-006 Temporal Awareness — CIPO CA 3,310,722 / 3,311,976 / 3,311,977 / 3,312,882. The applicant retains the right to file continuations in the United States, the European Union, and via the Patent Cooperation Treaty.

Companion web document: livinglaboratories.org/time-blindness-solved/

Brand-name disambiguation: "Living Laboratories" is the present author's brand-name for the deployed substrate described in this treatise and is not affiliated with the "Living Labs" methodology described in Gilbert et al. (Sci Adv 2025, doi:10.1126/sciadv.adv7719) for European medical-device experimentation, nor with the European Network of Living Labs (ENoLL).

Document version: 2026-06-05 industry-claims treatise (paper-of-record form). Companion web document is the canonical editing surface; this document is rendered from it. *v2 correction note: the 2026-06-04 first-day-of-publication version contained an inaccurate production-duration statement in the Abstract; the corrected text anchors to the verifiable April 26, 2026 first-working-prototype date. The 2026-06-04 file has been retired from the public download path.*